



US010803015B2

(12) **United States Patent**
Frerking et al.

(10) **Patent No.:** **US 10,803,015 B2**
(45) **Date of Patent:** ***Oct. 13, 2020**

(54) **CACHING SYSTEM AND METHOD**

G06F 12/0871 (2013.01); *G06F 16/23* (2019.01); *G06F 16/907* (2019.01); (Continued)

(71) Applicant: **Walmart Apollo, LLC**, Bentonville, AR (US)

(58) **Field of Classification Search**

CPC *G06F 16/172*; *G06F 16/907*; *G06F 16/23*; *G06F 3/0604*; *G06F 11/2023*; *G06F 11/2041*; *G06F 11/2097*; *G06F 12/0871*; *G06F 21/50*; *G06F 2201/82*;

(72) Inventors: **John Randy Frerking**, Prairie Grove, AR (US); **Richard Jackson**, Bella Vista, AR (US)

(Continued)

(73) Assignee: **WALMART APOLLO, LLC**, Bentonville, AR (US)

(56) **References Cited**

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 168 days.

U.S. PATENT DOCUMENTS

This patent is subject to a terminal disclaimer.

6,330,606 B1 12/2001 Logue et al.
6,611,498 B1 8/2003 Baker et al.
(Continued)

(21) Appl. No.: **15/953,949**

OTHER PUBLICATIONS

(22) Filed: **Apr. 16, 2018**

International Search Report & Written Opinion in related International Patent Application No. PCT/US14/49976, dated Nov. 6, 2014; 9 pages.

(65) **Prior Publication Data**

US 2018/0246903 A1 Aug. 30, 2018

(Continued)

Related U.S. Application Data

Primary Examiner — Hanh B Thai

(63) Continuation-in-part of application No. 15/143,953, filed on May 2, 2016, now Pat. No. 10,419,572, and (Continued)

(74) *Attorney, Agent, or Firm* — Schmeiser, Olsen & Watts LLP

(51) **Int. Cl.**
G06F 16/00 (2019.01)
G06F 16/172 (2019.01)

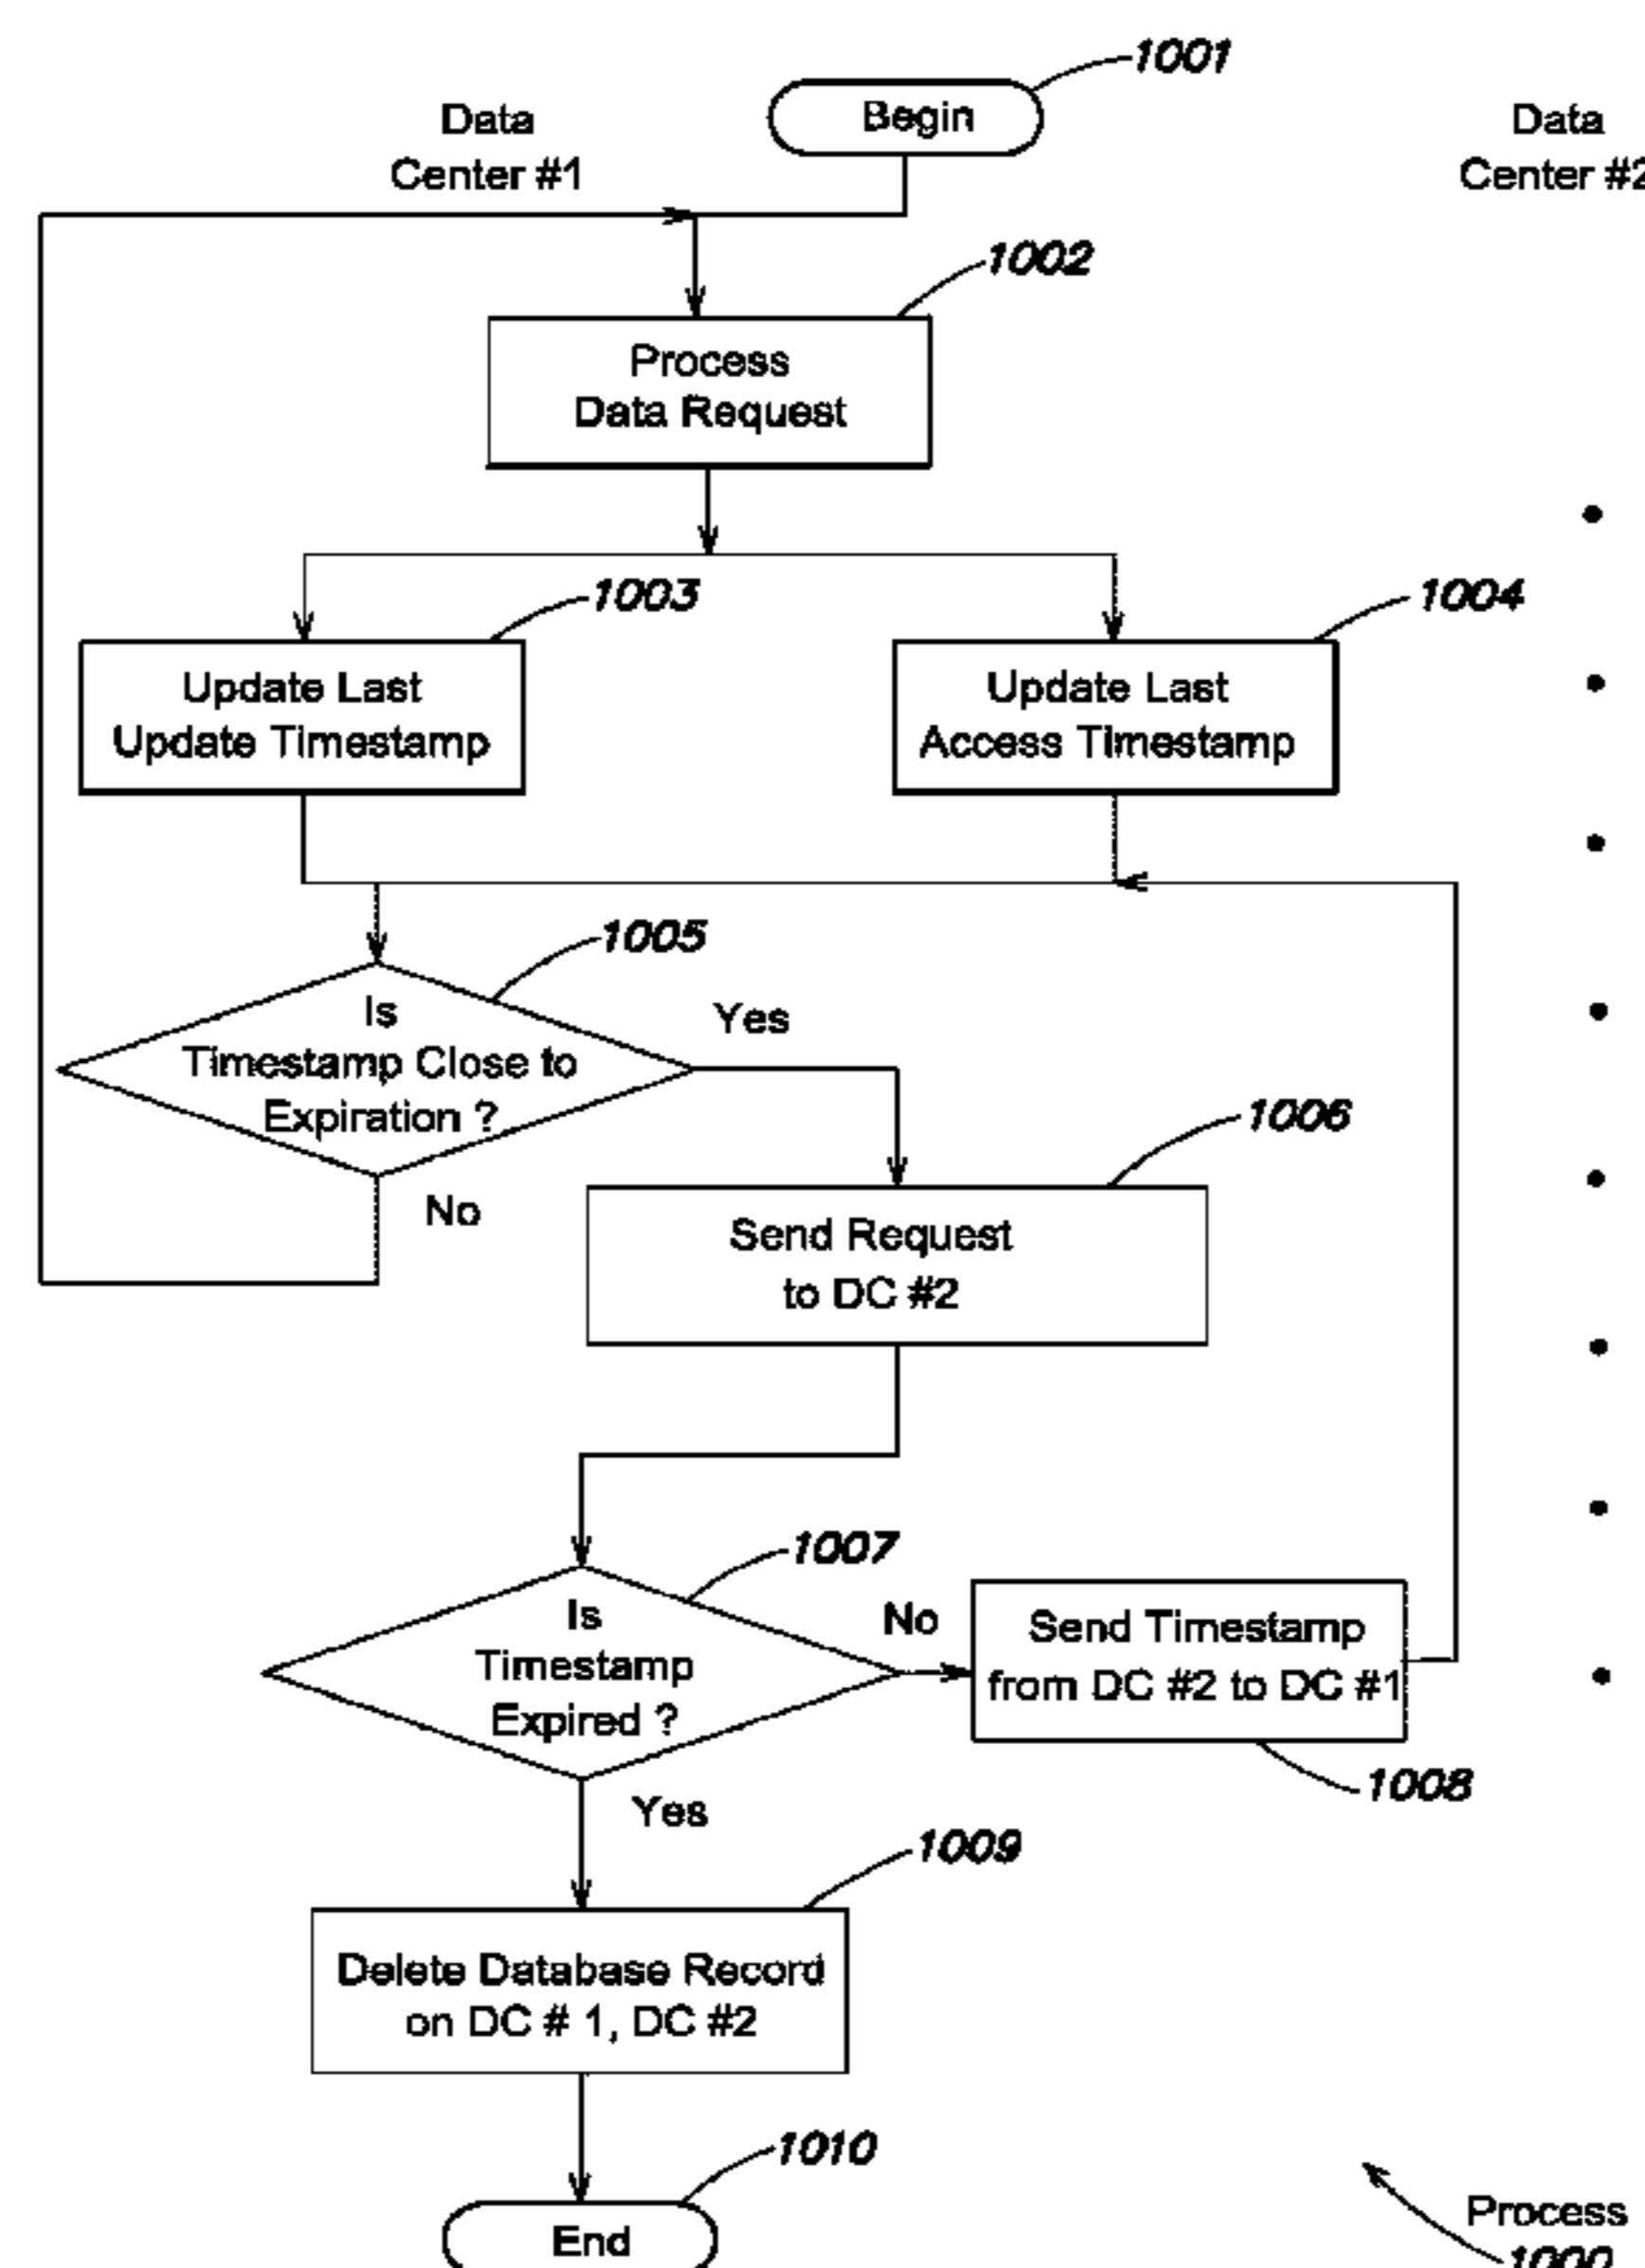
(Continued)

(57) **ABSTRACT**

Disclosed is a caching system for processing web-based service requests. The caching system includes one or more pluralities of servers coupled to one or more online common storage. Each one of the plurality of servers processes caching service requests from one or more clients coupled to the plurality of servers. Each one of the plurality of servers is capable of operating in an active/single mode, and active/active mode, or an active/standby mode. A timestamp exchange process is used to update or eliminate outdated database records.

(52) **U.S. Cl.**
CPC *G06F 16/172* (2019.01); *G06F 3/0604* (2013.01); *G06F 11/2023* (2013.01); *G06F 11/2041* (2013.01); *G06F 11/2097* (2013.01);

7 Claims, 25 Drawing Sheets



Related U.S. Application Data

a continuation-in-part of application No. 13/960,192, filed on Aug. 6, 2013, now Pat. No. 10,116,762, application No. 15/143,953, which is a continuation-in-part of application No. 13/960,266, filed on Aug. 6, 2013, now Pat. No. 9,336,265.

(51) **Int. Cl.**

G06F 12/0871 (2016.01)
G06F 21/50 (2013.01)
H04L 29/06 (2006.01)
H04L 29/08 (2006.01)
G06F 3/06 (2006.01)
G06F 16/23 (2019.01)
G06F 16/907 (2019.01)
G06F 11/20 (2006.01)

(52) **U.S. Cl.**

CPC **G06F 21/50** (2013.01); **H04L 63/10** (2013.01); **H04L 67/1004** (2013.01); **H04L 67/2842** (2013.01); **H04L 67/42** (2013.01); **G06F 2201/82** (2013.01); **G06F 2201/835** (2013.01)

(58) **Field of Classification Search**

CPC G06F 2201/835; H04L 63/10; H04L 67/1004; H04L 67/2842; H04L 67/42
 USPC 707/638
 See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,629,144 B1 * 9/2003 Chu G06F 11/1443 709/227
 6,704,873 B1 3/2004 Underwood
 7,139,973 B1 11/2006 Kirkwood et al.
 7,165,116 B2 1/2007 Grove et al.
 7,349,871 B2 3/2008 Labrou et al.
 7,558,854 B2 7/2009 Nakahara et al.
 7,801,826 B2 9/2010 Labrou et al.
 7,865,584 B2 1/2011 Grossner et al.
 7,962,603 B1 6/2011 Morimoto
 8,219,676 B2 7/2012 Jagadish et al.
 8,271,430 B2 9/2012 Wilson
 8,938,636 B1 1/2015 Hochschild et al.
 9,055,139 B1 6/2015 Devireddy
 9,336,265 B2 5/2016 Frerking
 9,529,772 B1 * 12/2016 Shankaran H04L 41/0816
 9,602,614 B1 * 3/2017 Shankaran G06F 12/0806
 2002/0085579 A1 7/2002 Sullivan
 2002/0116474 A1 8/2002 Copeland et al.
 2002/0116582 A1 8/2002 Copeland et al.
 2003/0041227 A1 2/2003 Nakamatsu
 2003/0065941 A1 4/2003 Ballard et al.
 2004/0098350 A1 5/2004 Labrou et al.
 2004/0107170 A1 6/2004 Labrou et al.
 2004/0205099 A1 10/2004 Hagiwara
 2004/0213387 A1 10/2004 Chandrasekaran
 2005/0027543 A1 2/2005 Labrou et al.
 2005/0028024 A1 2/2005 Kataoka et al.
 2005/0044197 A1 2/2005 Lai

2005/0050112 A1 * 3/2005 Chandrasekaran G06F 16/2322
 2005/0198200 A1 9/2005 Subramanian et al.
 2006/0165040 A1 7/2006 Rathod
 2006/0195607 A1 8/2006 Naseh et al.
 2007/0203944 A1 8/2007 Baltra et al.
 2007/0282858 A1 12/2007 Arner
 2008/0082623 A1 4/2008 Michael et al.
 2008/0140640 A1 * 6/2008 Raff H04L 29/06
 2008/0235298 A1 9/2008 Lin et al.
 2009/0043881 A1 2/2009 Alstad
 2009/0113531 A1 4/2009 Emmerich et al.
 2009/0287746 A1 11/2009 Brown
 2009/0299987 A1 12/2009 Wilson
 2009/0327098 A1 12/2009 Ronen et al.
 2010/0138485 A1 * 6/2010 Chow H04L 67/2842 709/203
 2010/0198972 A1 8/2010 Umbehooker
 2010/0299553 A1 11/2010 Cen
 2010/0325371 A1 12/2010 Jagadish et al.
 2011/0038633 A1 2/2011 DeCusatis et al.
 2011/0072217 A1 3/2011 Hoang et al.
 2011/0131197 A1 6/2011 Nielsen et al.
 2011/0138027 A1 6/2011 Friedmann et al.
 2011/0208695 A1 8/2011 Anand et al.
 2012/0226712 A1 9/2012 Vermeulen et al.
 2013/0031060 A1 * 1/2013 Lowery G06F 16/9574 707/689
 2013/0174223 A1 4/2013 Dykeman et al.
 2013/0185716 A1 7/2013 Yin et al.
 2013/0304843 A1 11/2013 Chow et al.
 2013/0318191 A1 11/2013 Yin et al.
 2014/0067852 A1 3/2014 Wong et al.
 2014/0149537 A1 * 5/2014 Shankaran H04L 41/0816 709/216
 2014/0211793 A1 7/2014 Mital et al.
 2014/0359043 A1 12/2014 Gao et al.
 2015/0046511 A1 2/2015 Frerking
 2016/0248875 A1 8/2016 Frerking et al.

OTHER PUBLICATIONS

International Search Report & Written Opinion in related International Patent Application No. PCT/US14/49966, dated Nov. 6, 2014; 7 pages.
 Prosecution history of related U.S. Appl. No. 13/960,266, filed Aug. 6, 2013.
 Prosecution history of related U.S. Appl. No. 13/960,192, filed Aug. 6, 2013.
 Prosecution history of related U.S. Appl. No. 15/143,953, filed May 2, 2016.
 International Preliminary Report in related International Patent Application No. PCT/US14/49966, dated Feb. 18, 2016; 6 pages.
 International Preliminary Report in related International Patent Application No. PCT/US14/49976, dated Feb. 18, 2016; 8 pages.
 Non-Final Office Action in U.S. Appl. No. 15/143,953 dated Jan. 18, 2019; 9 pages.
 Notice of Allowance in U.S. Appl. No. 13/960,192 dated Aug. 14, 2018; 10 pages.
 Restriction Requirement in U.S. Appl. No. 15/143,953 dated Sep. 26, 2018; 5 pages.
 Notice of Allowance in U.S. Appl. No. 15/143,953 dated Jun. 12, 2019; 11 pages.

* cited by examiner

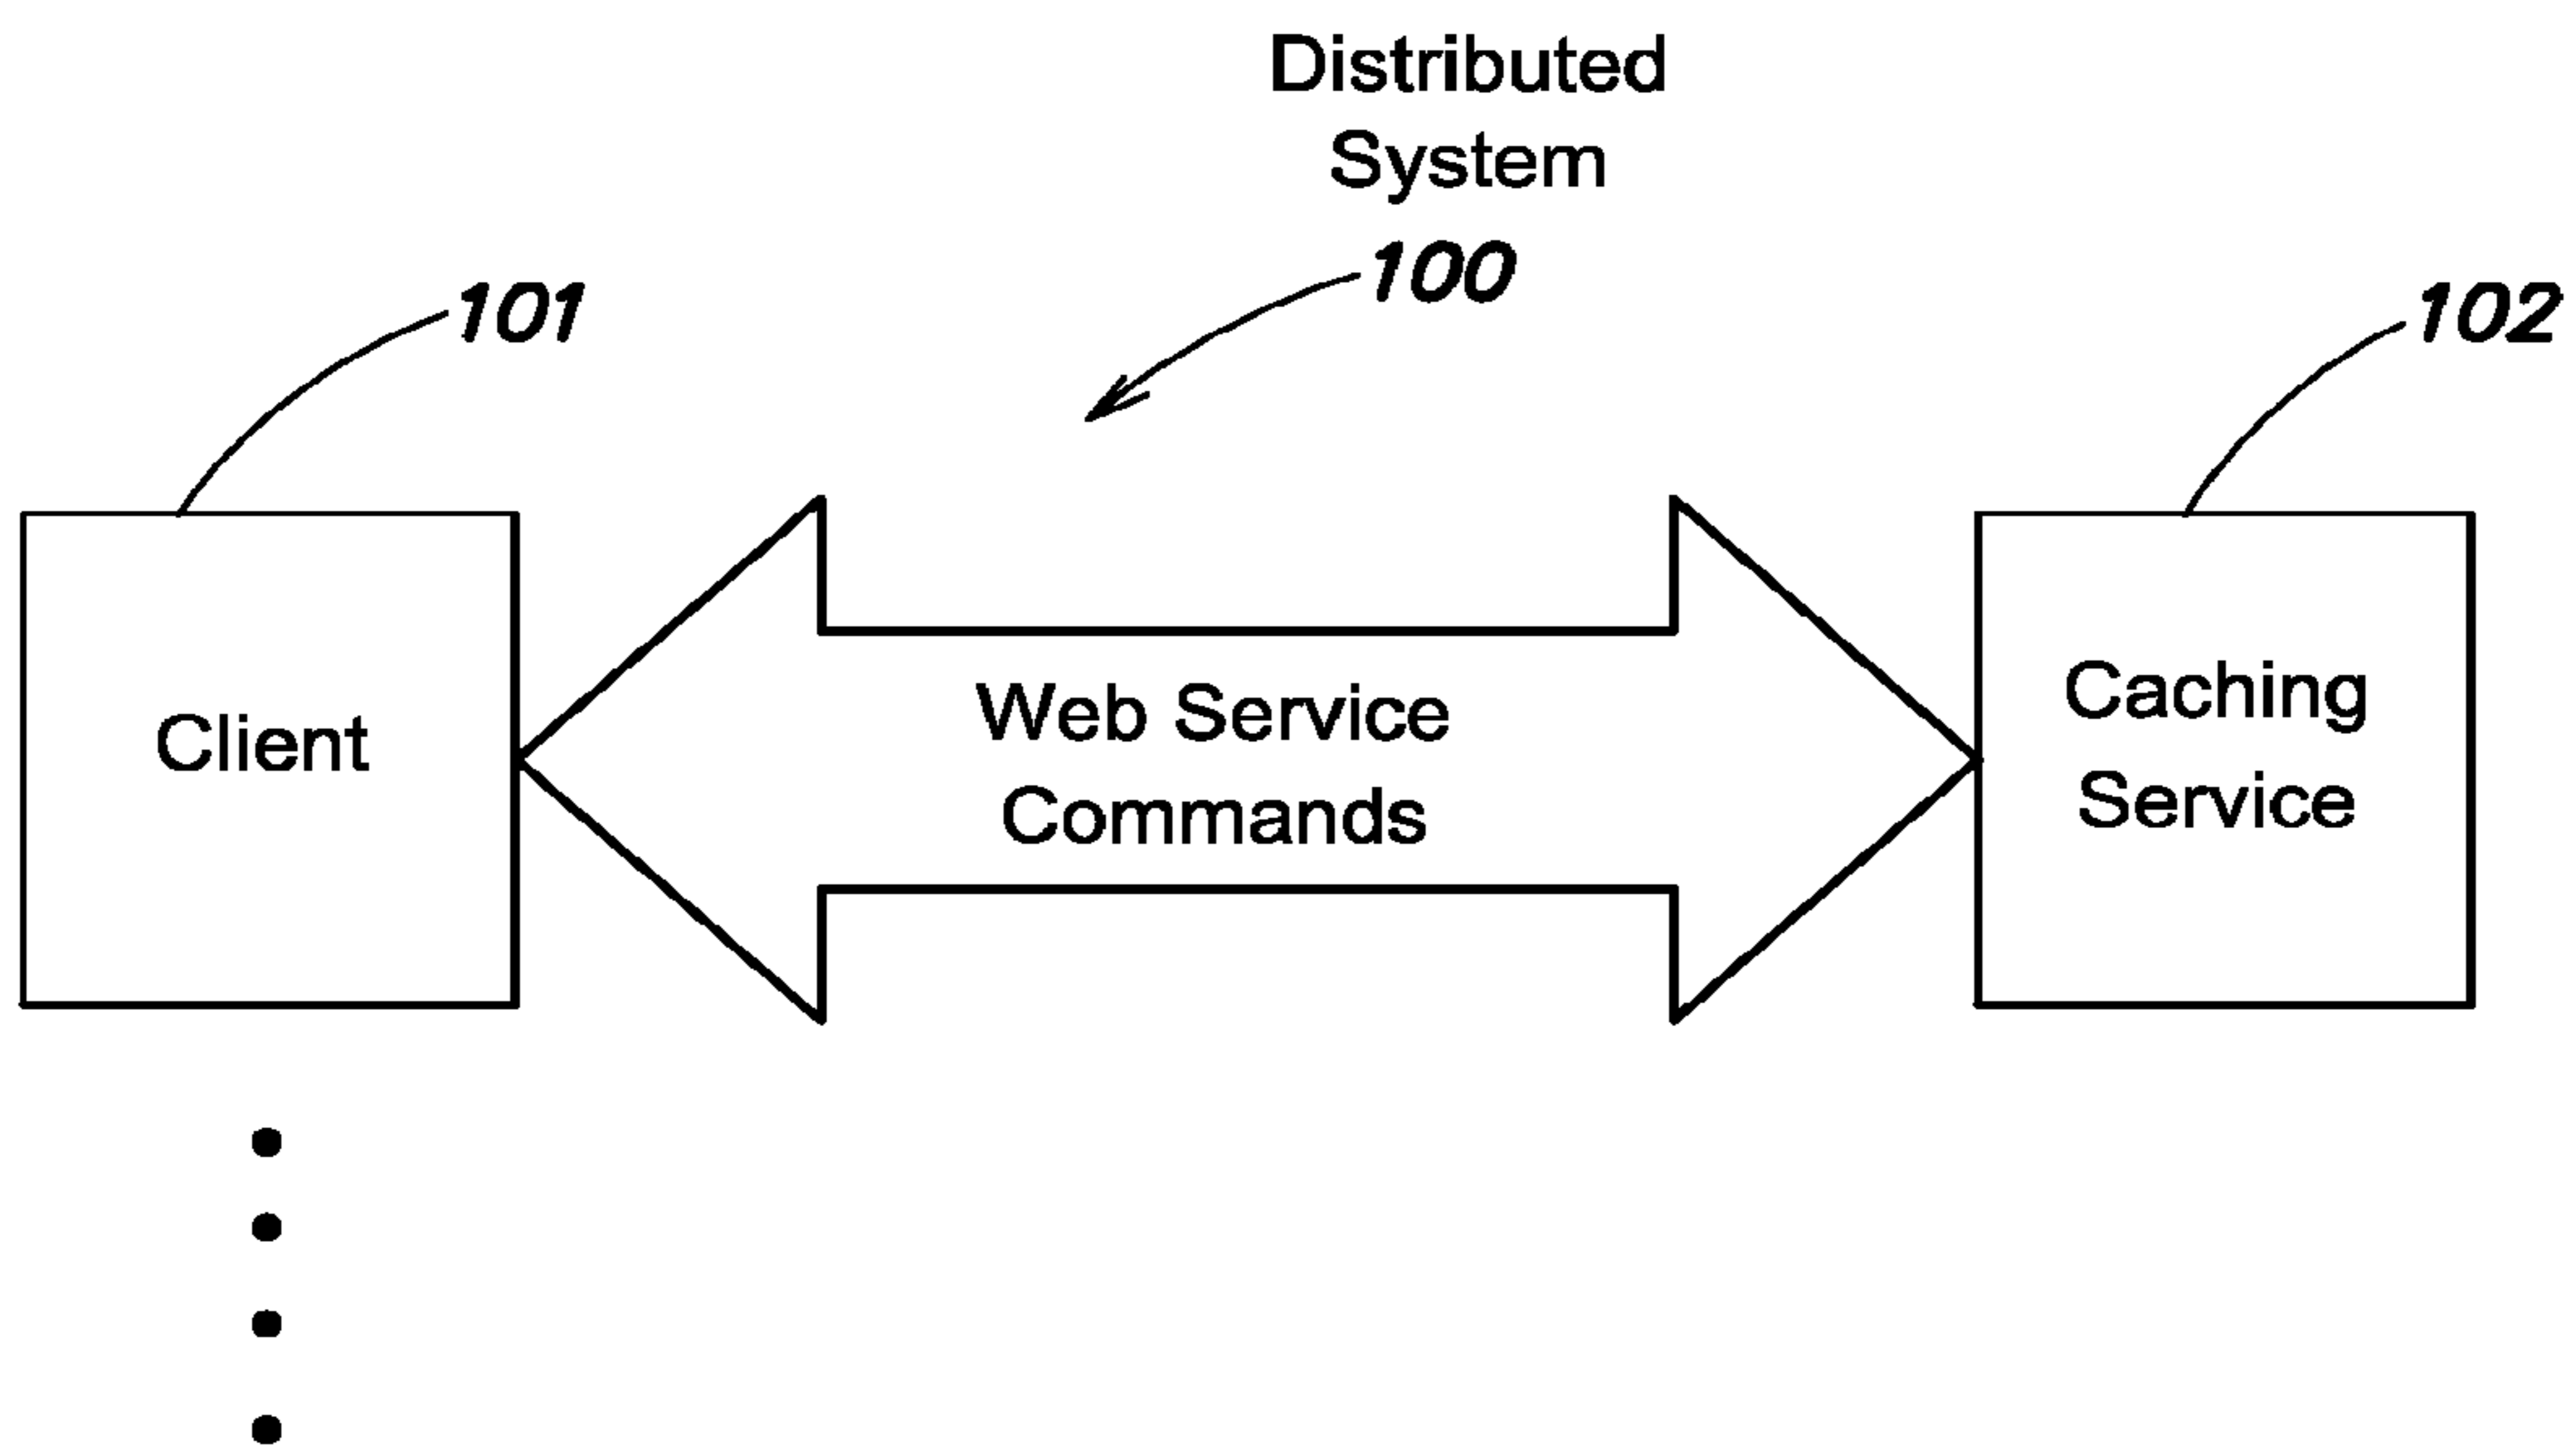


FIG. 1




FIG. 2A

FIG. 2B

FIG. 2C

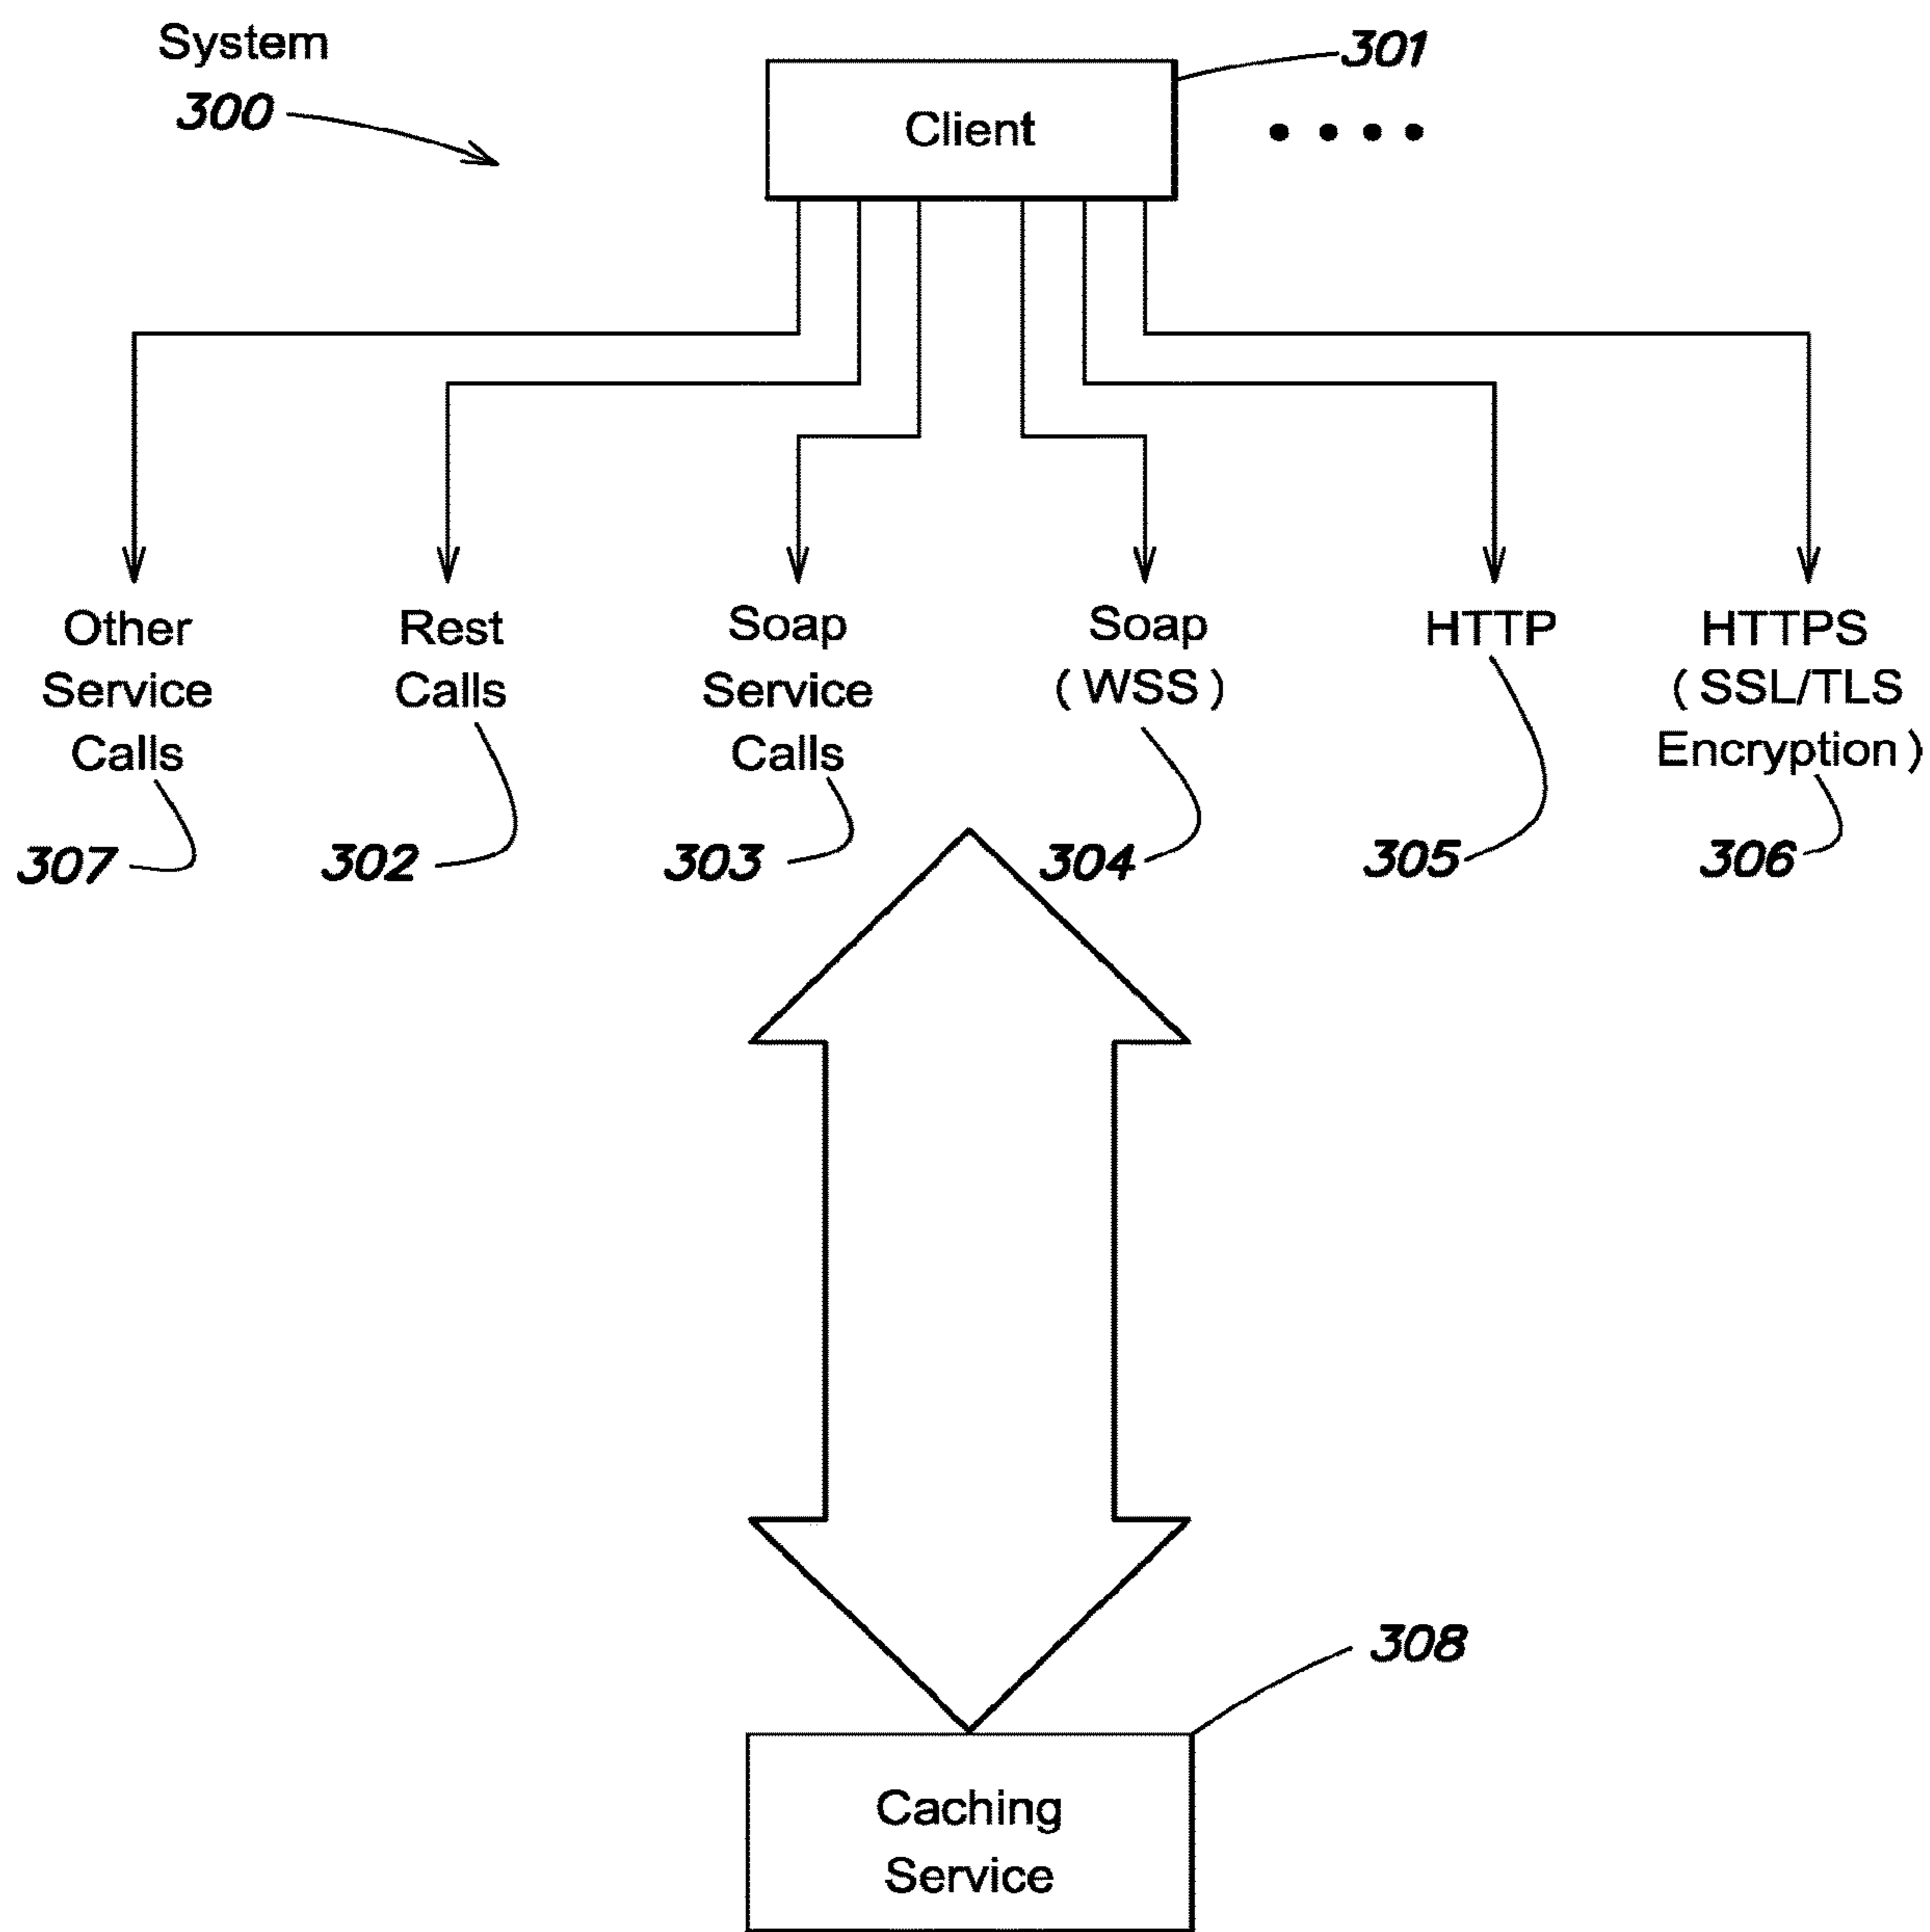


FIG. 3




FIG. 4

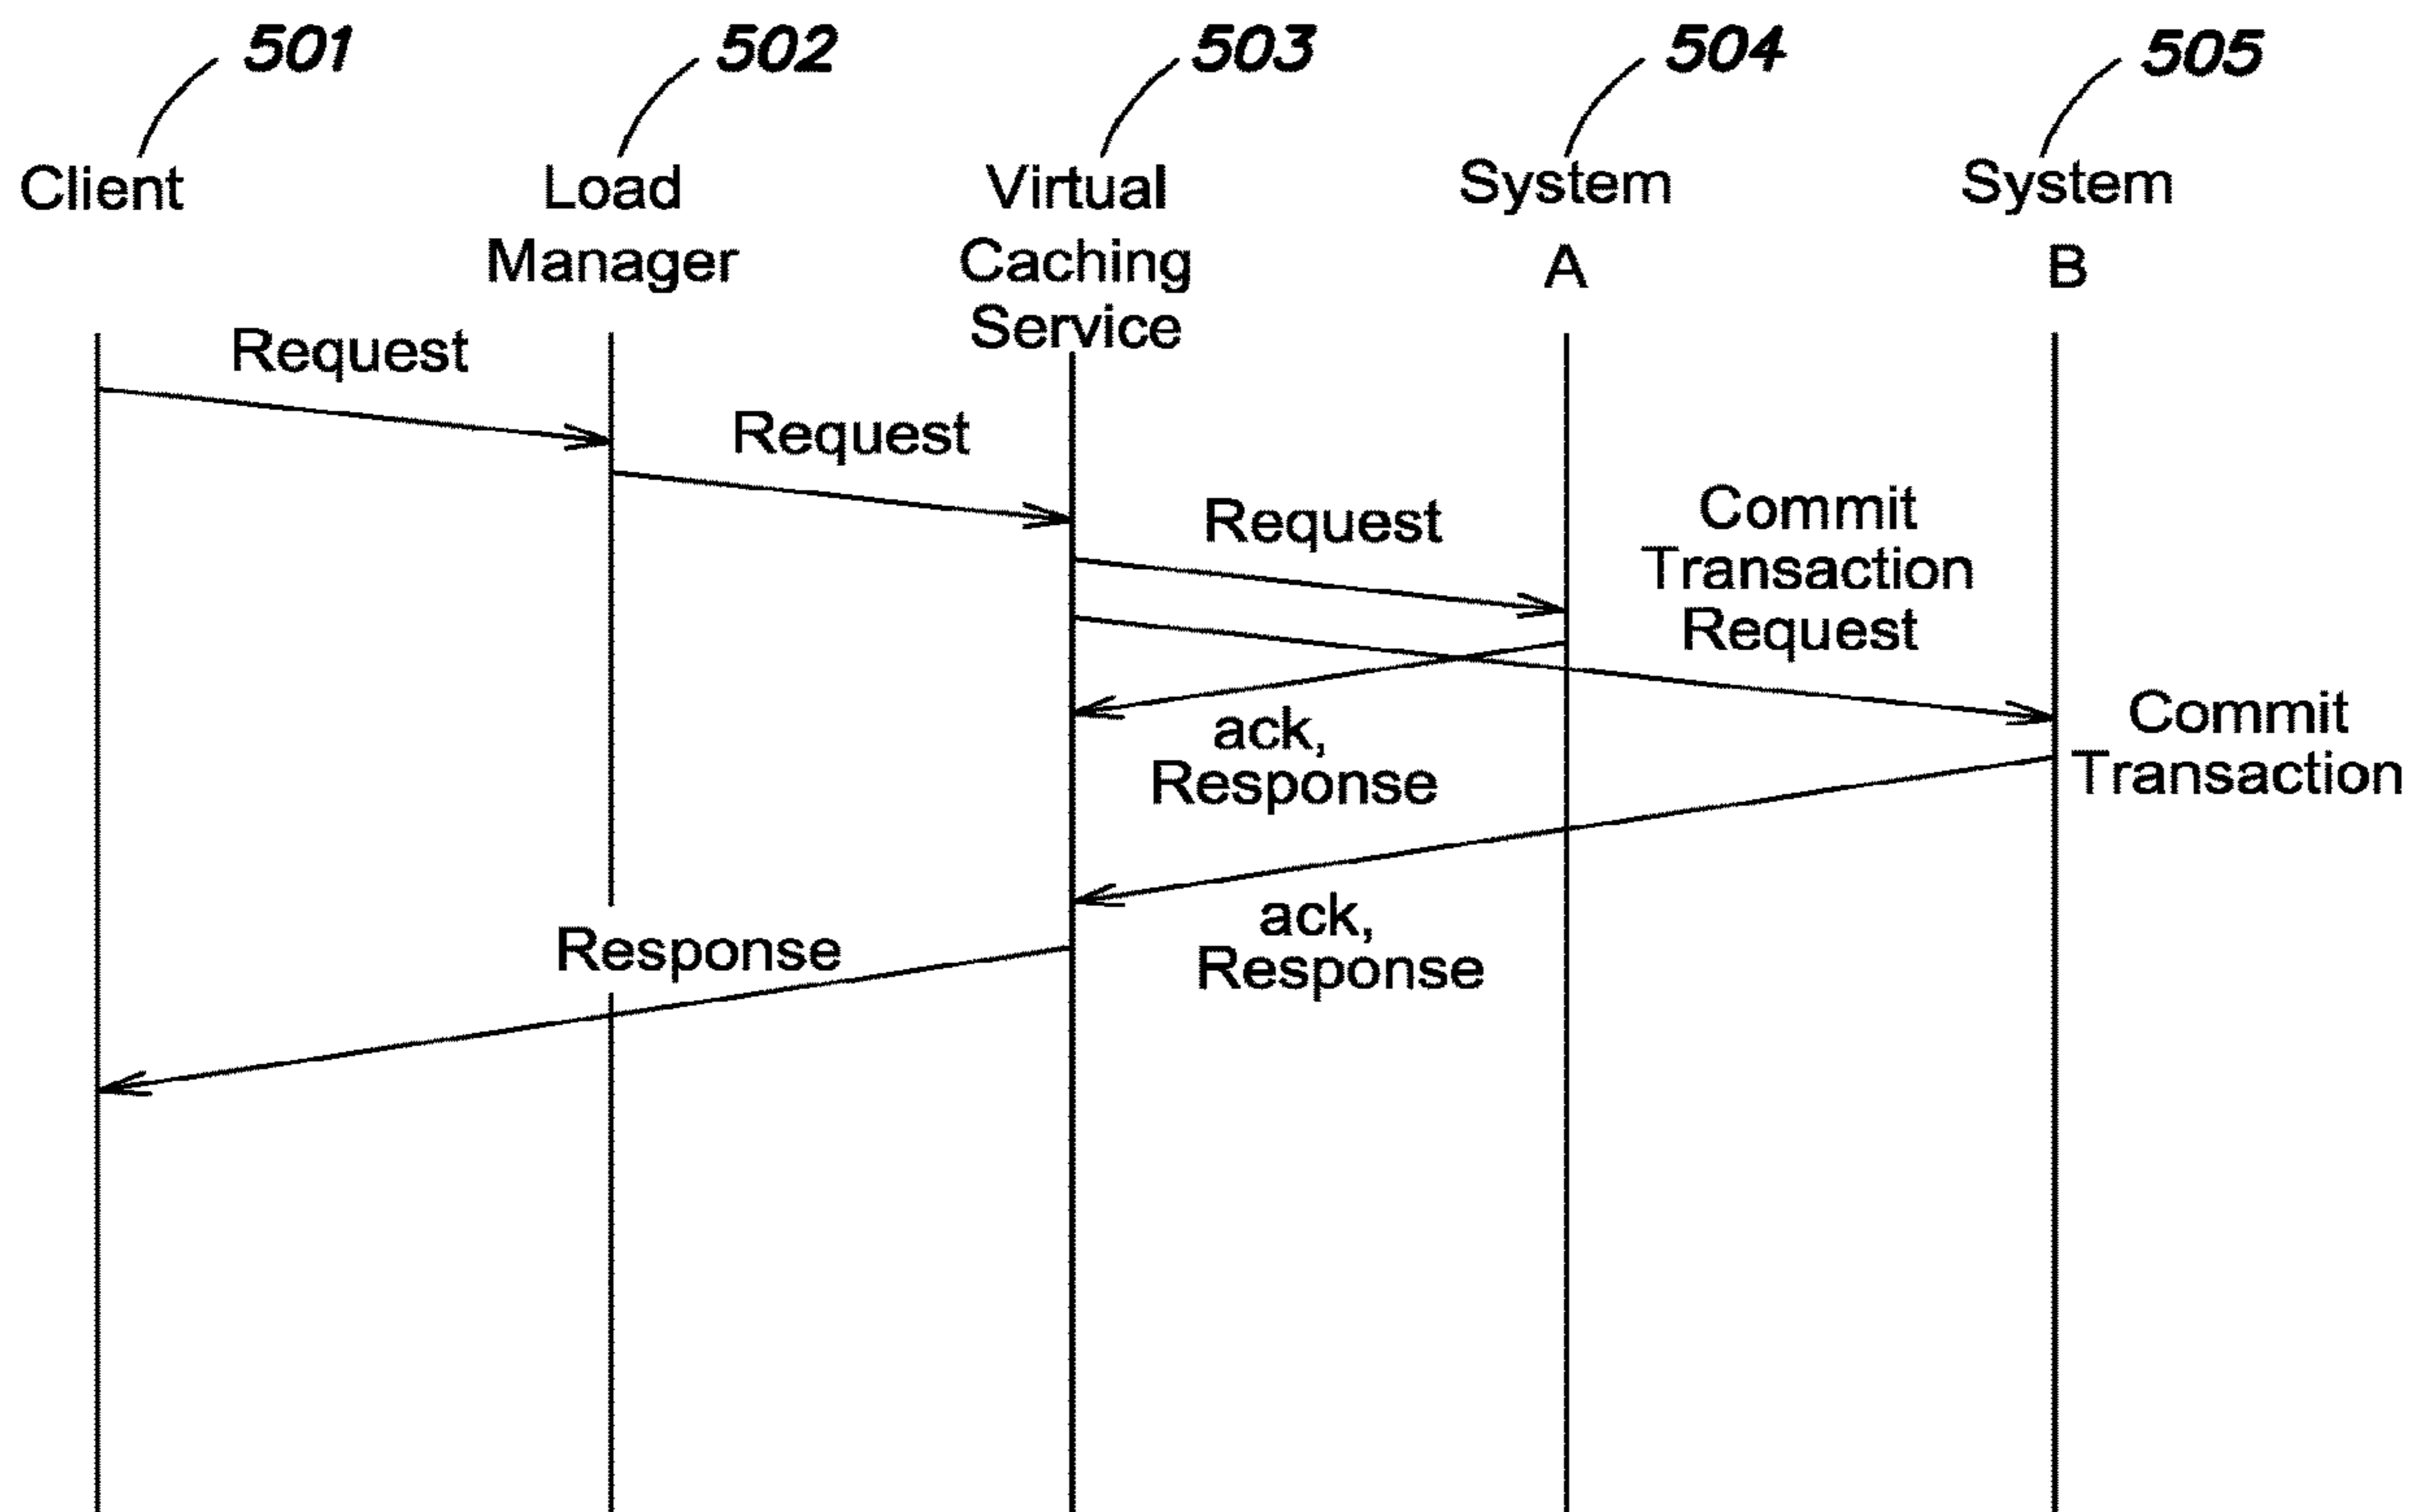


FIG. 5A

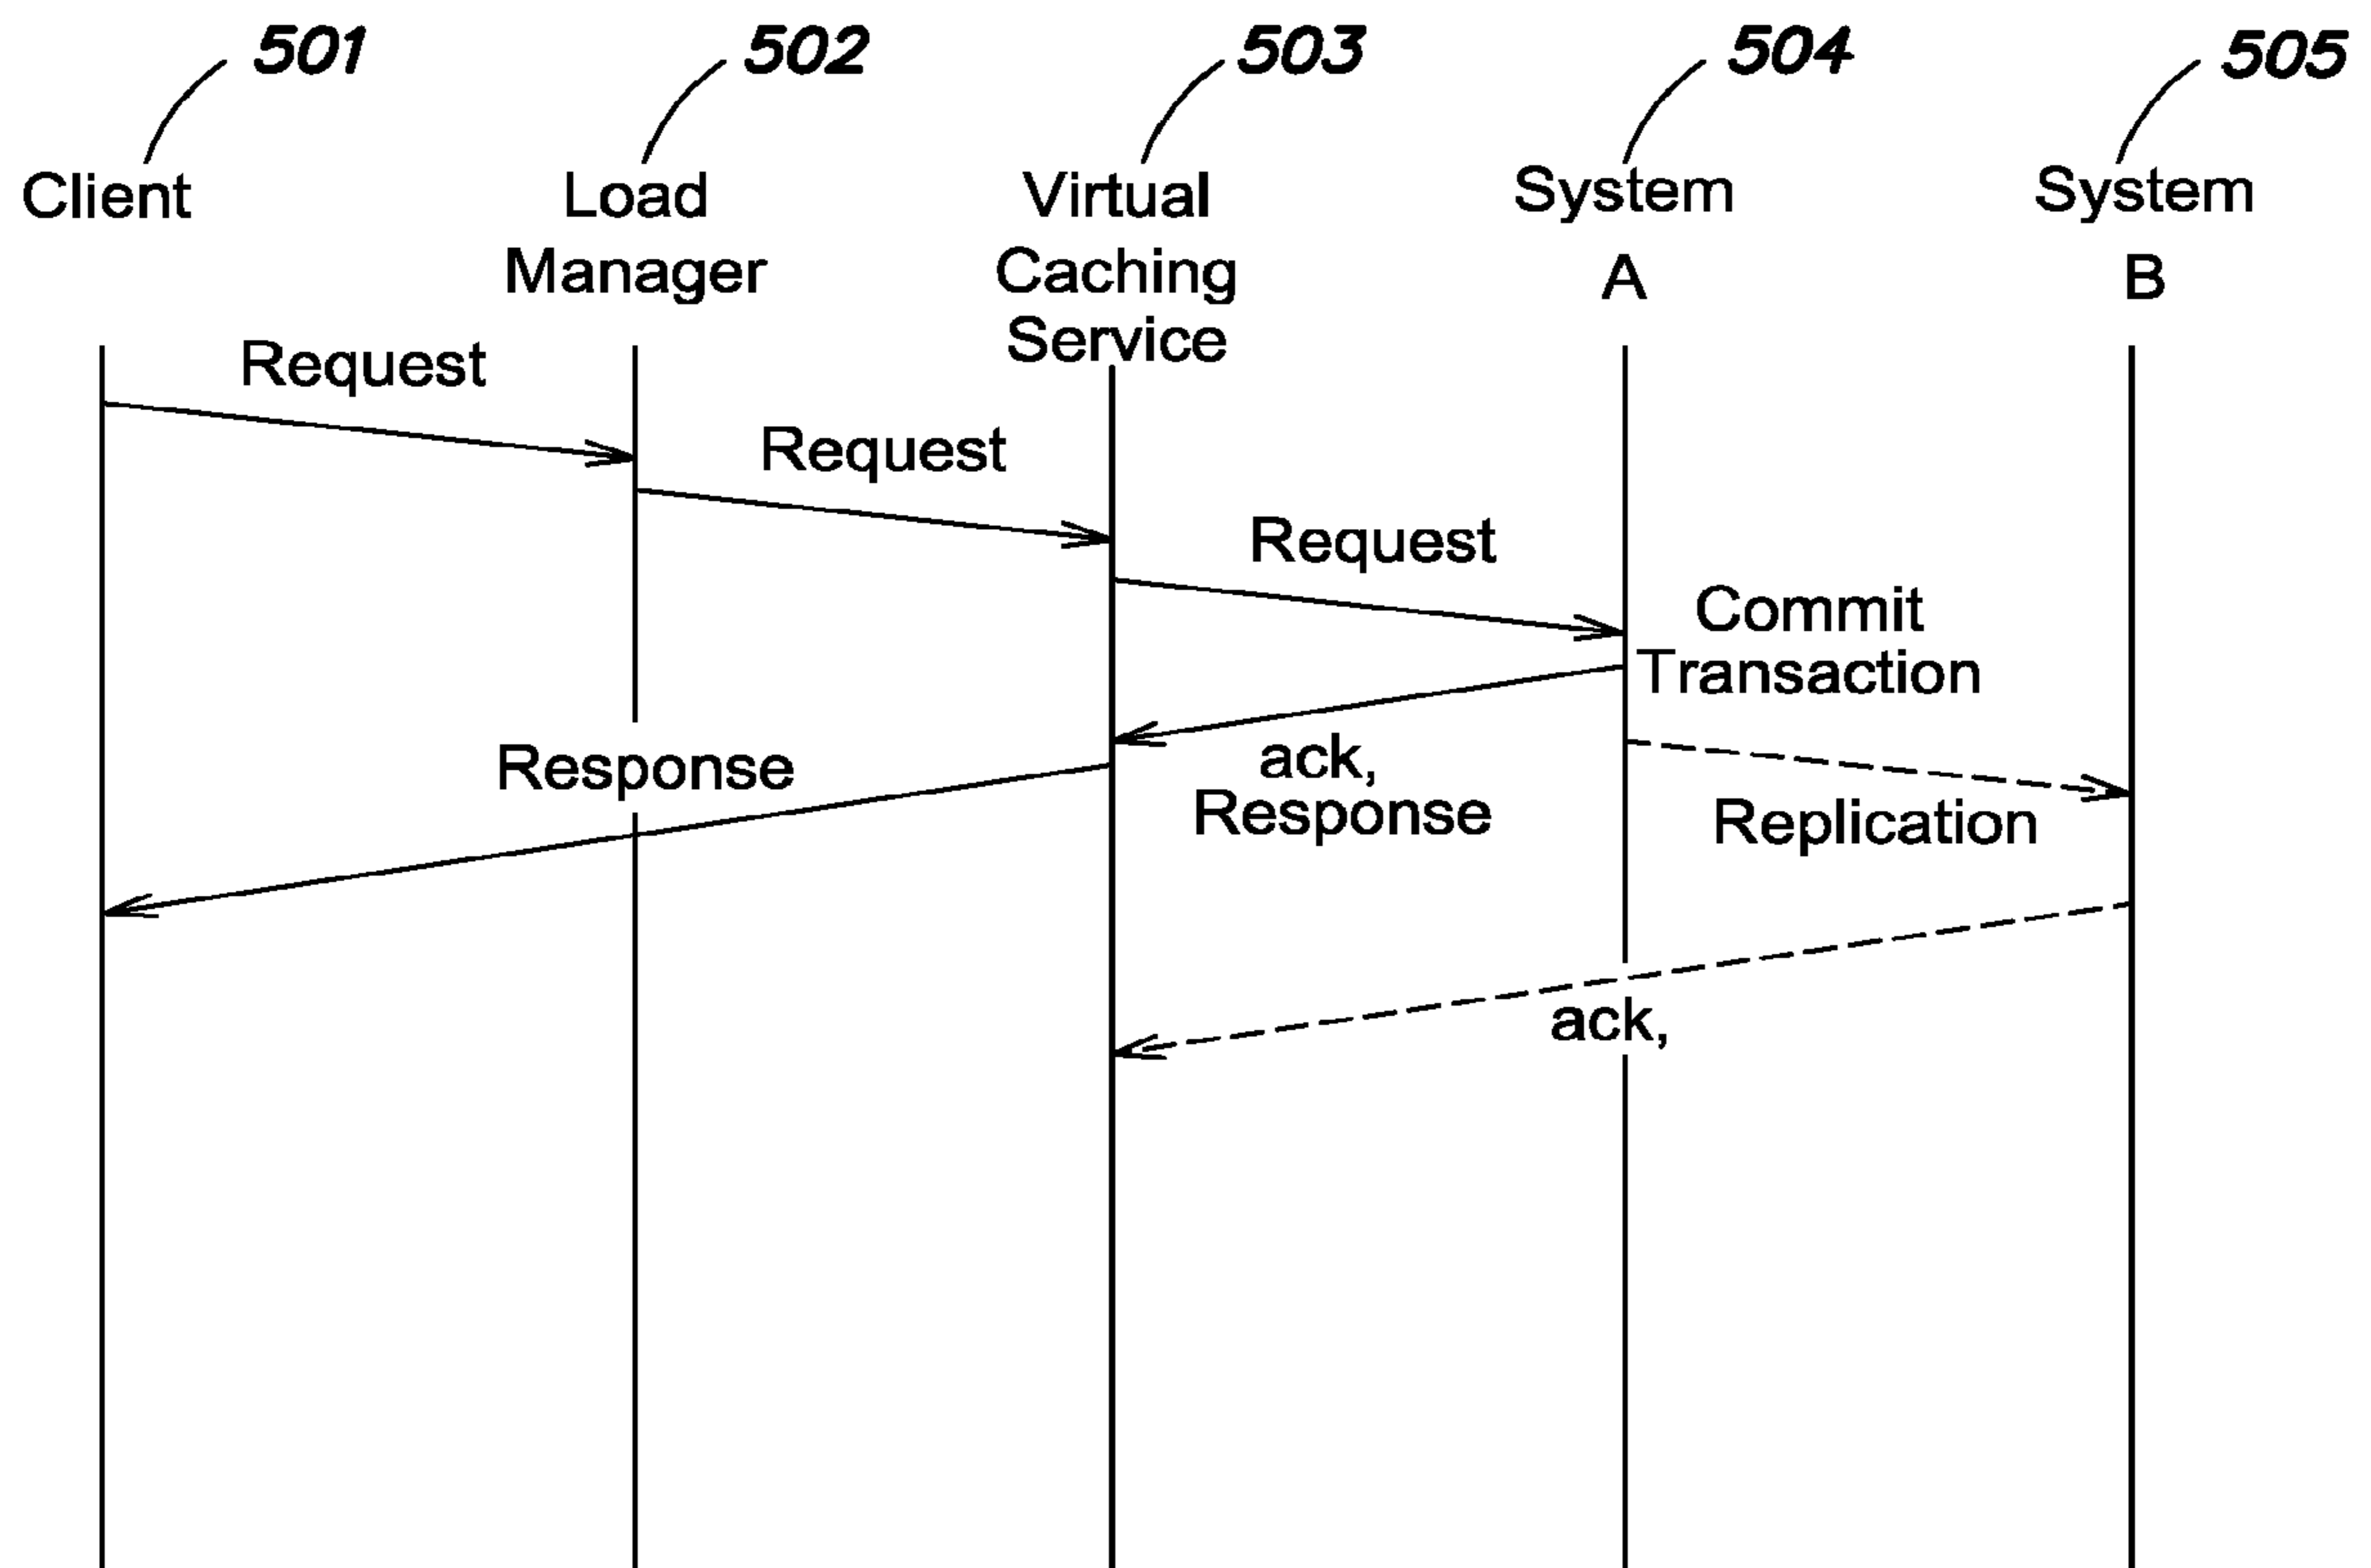


FIG. 5B




FIG. 6




FIG. 7




FIG. 8

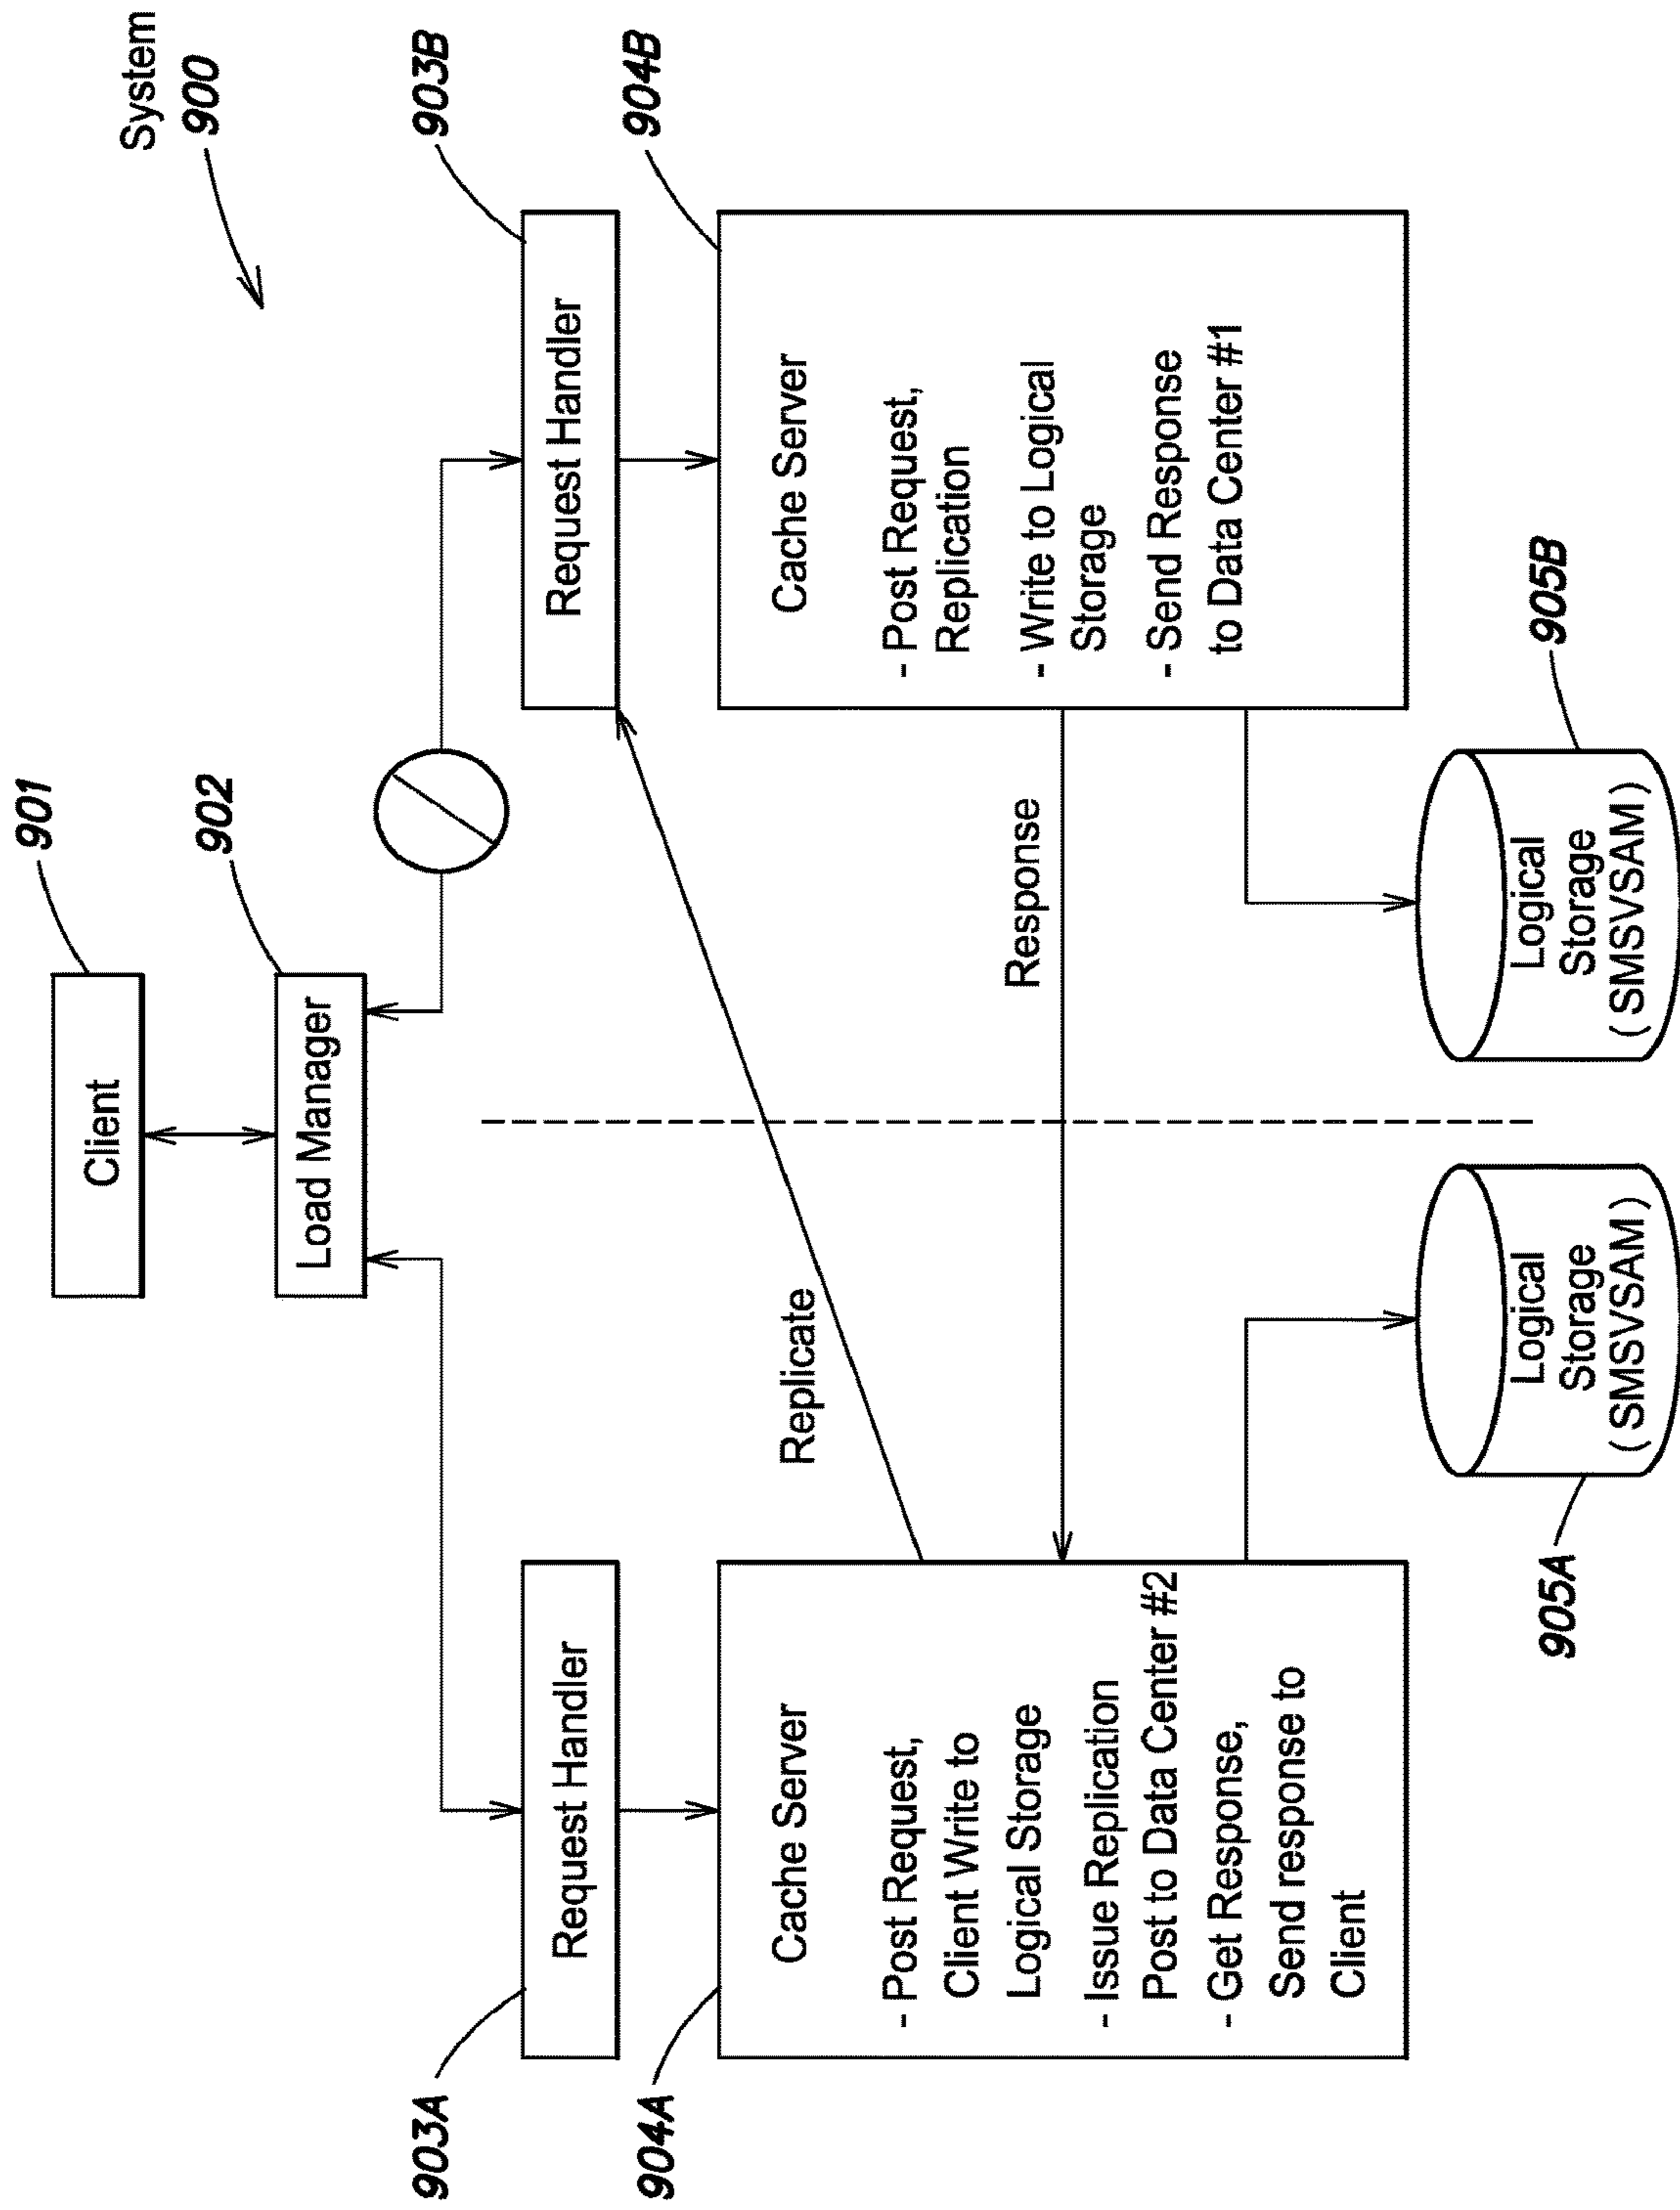


FIG. 9

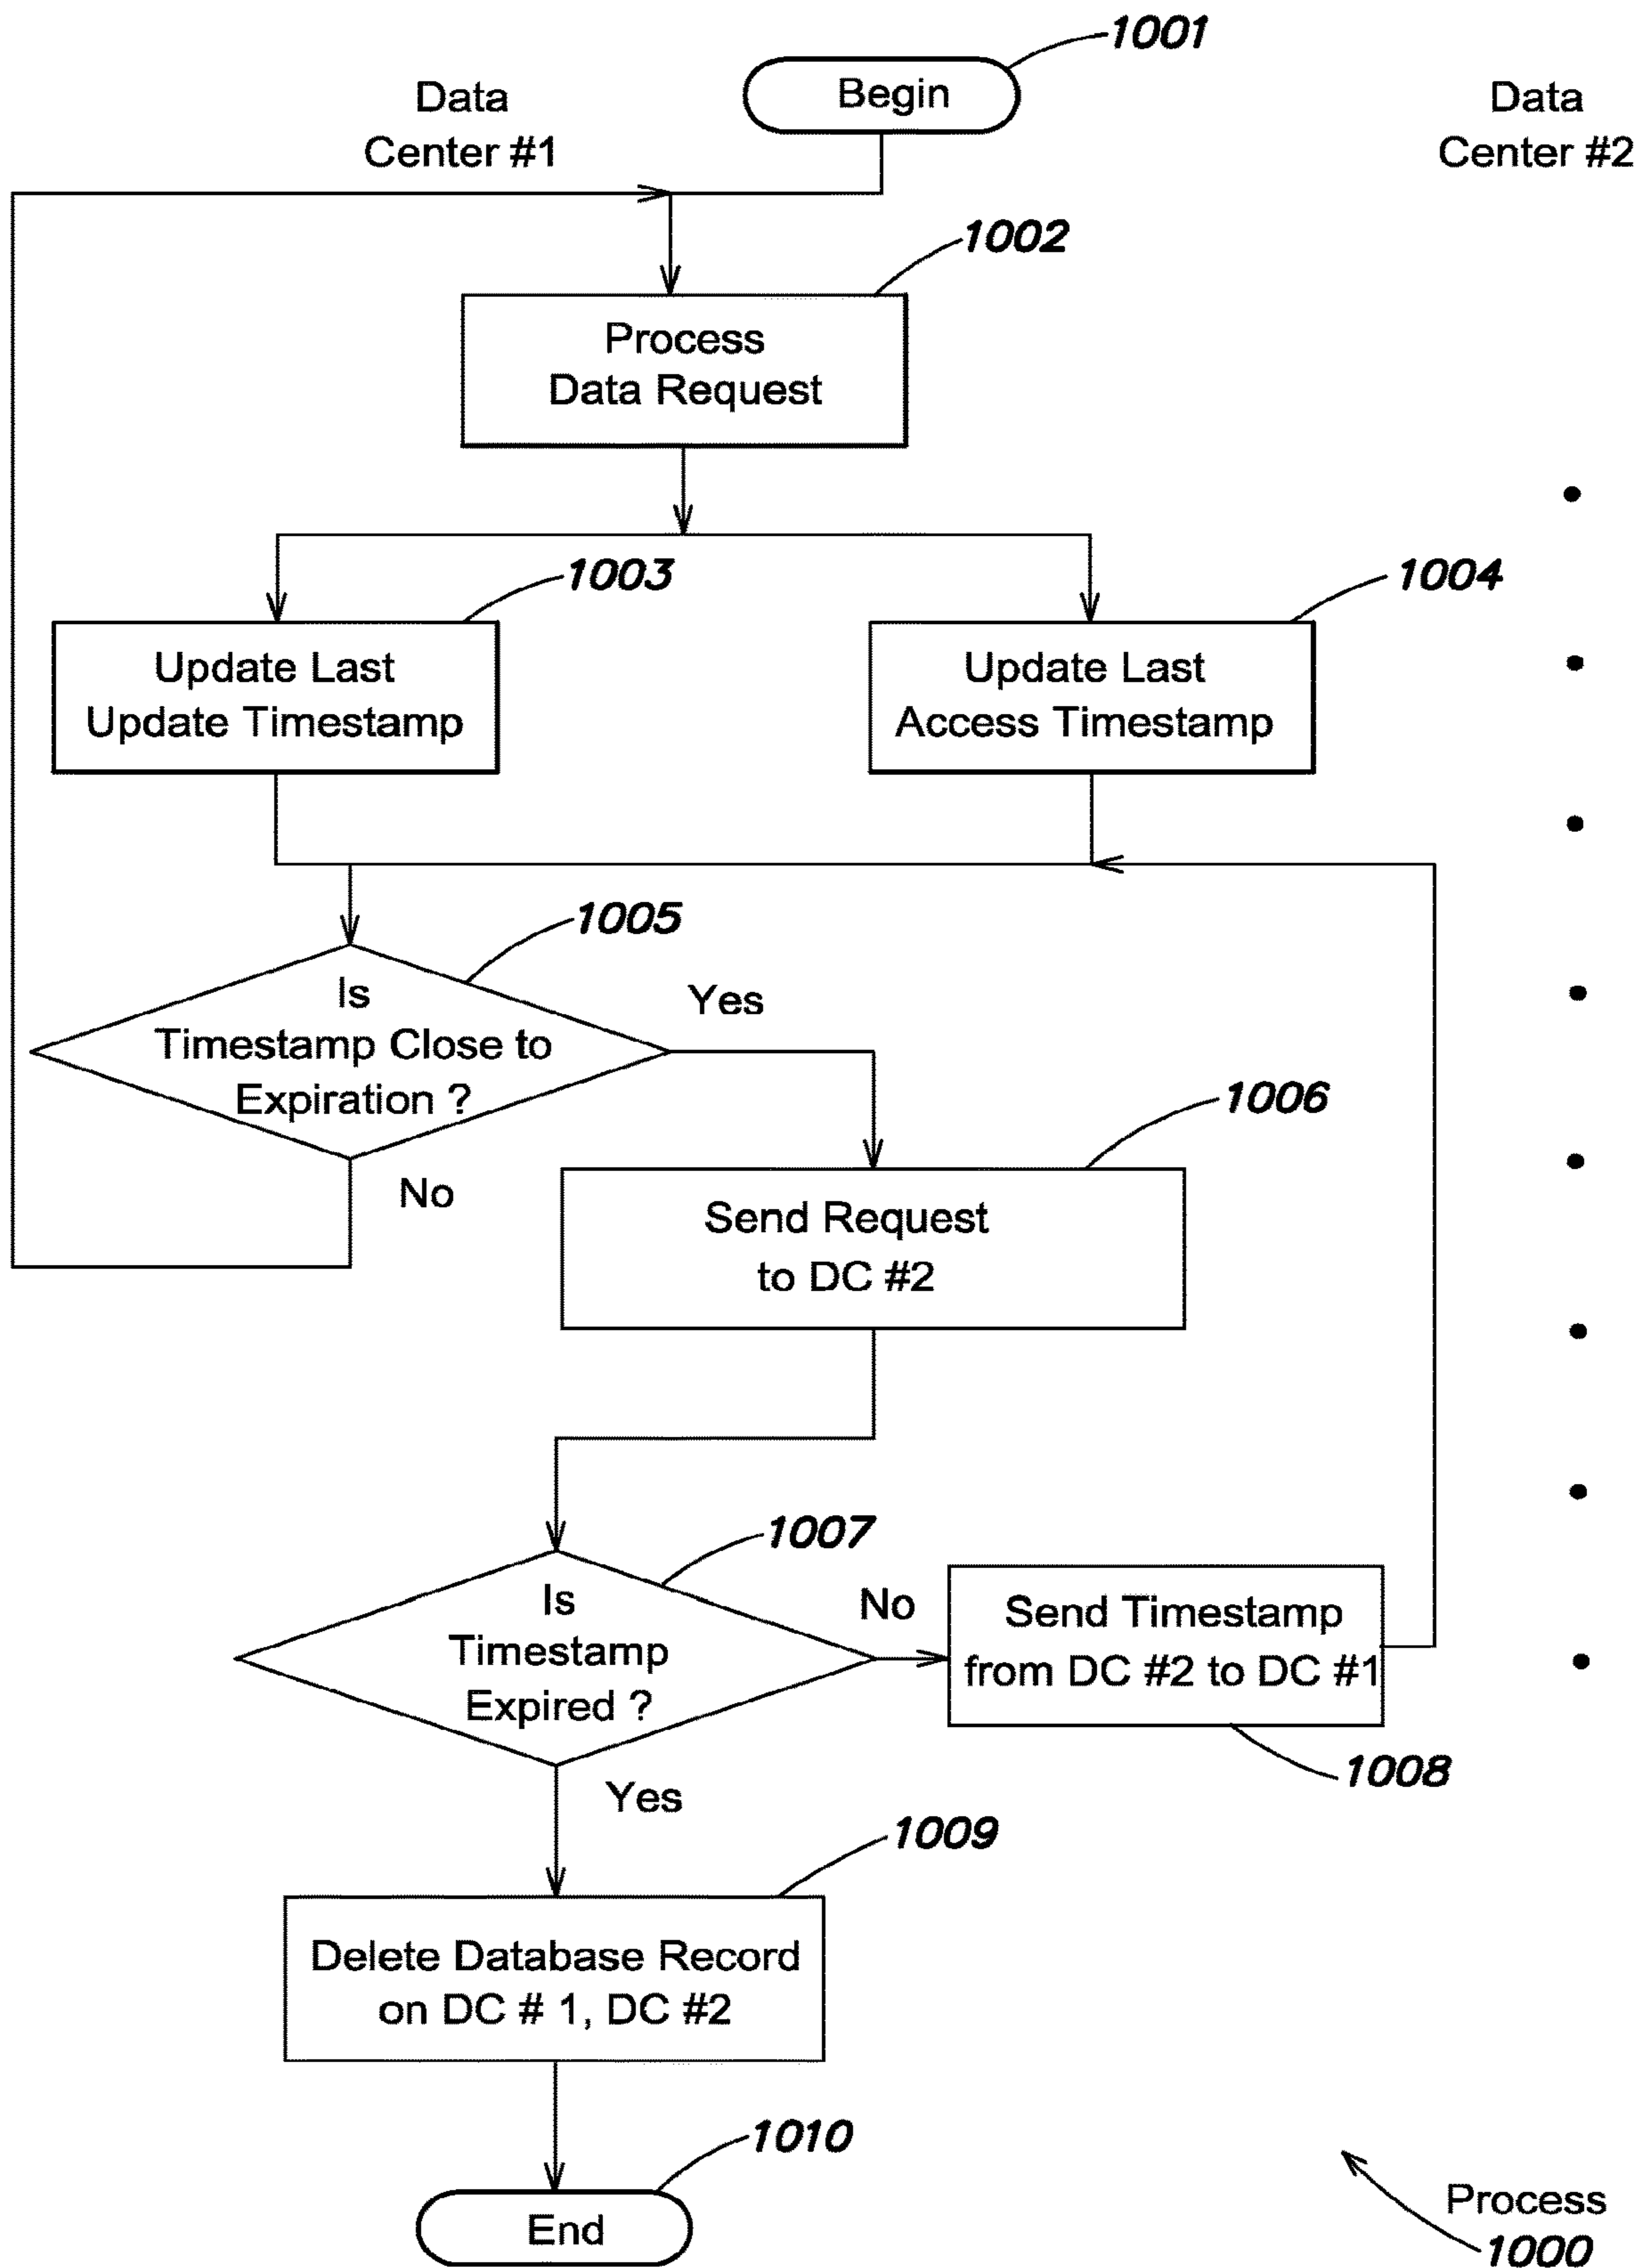


FIG. 10

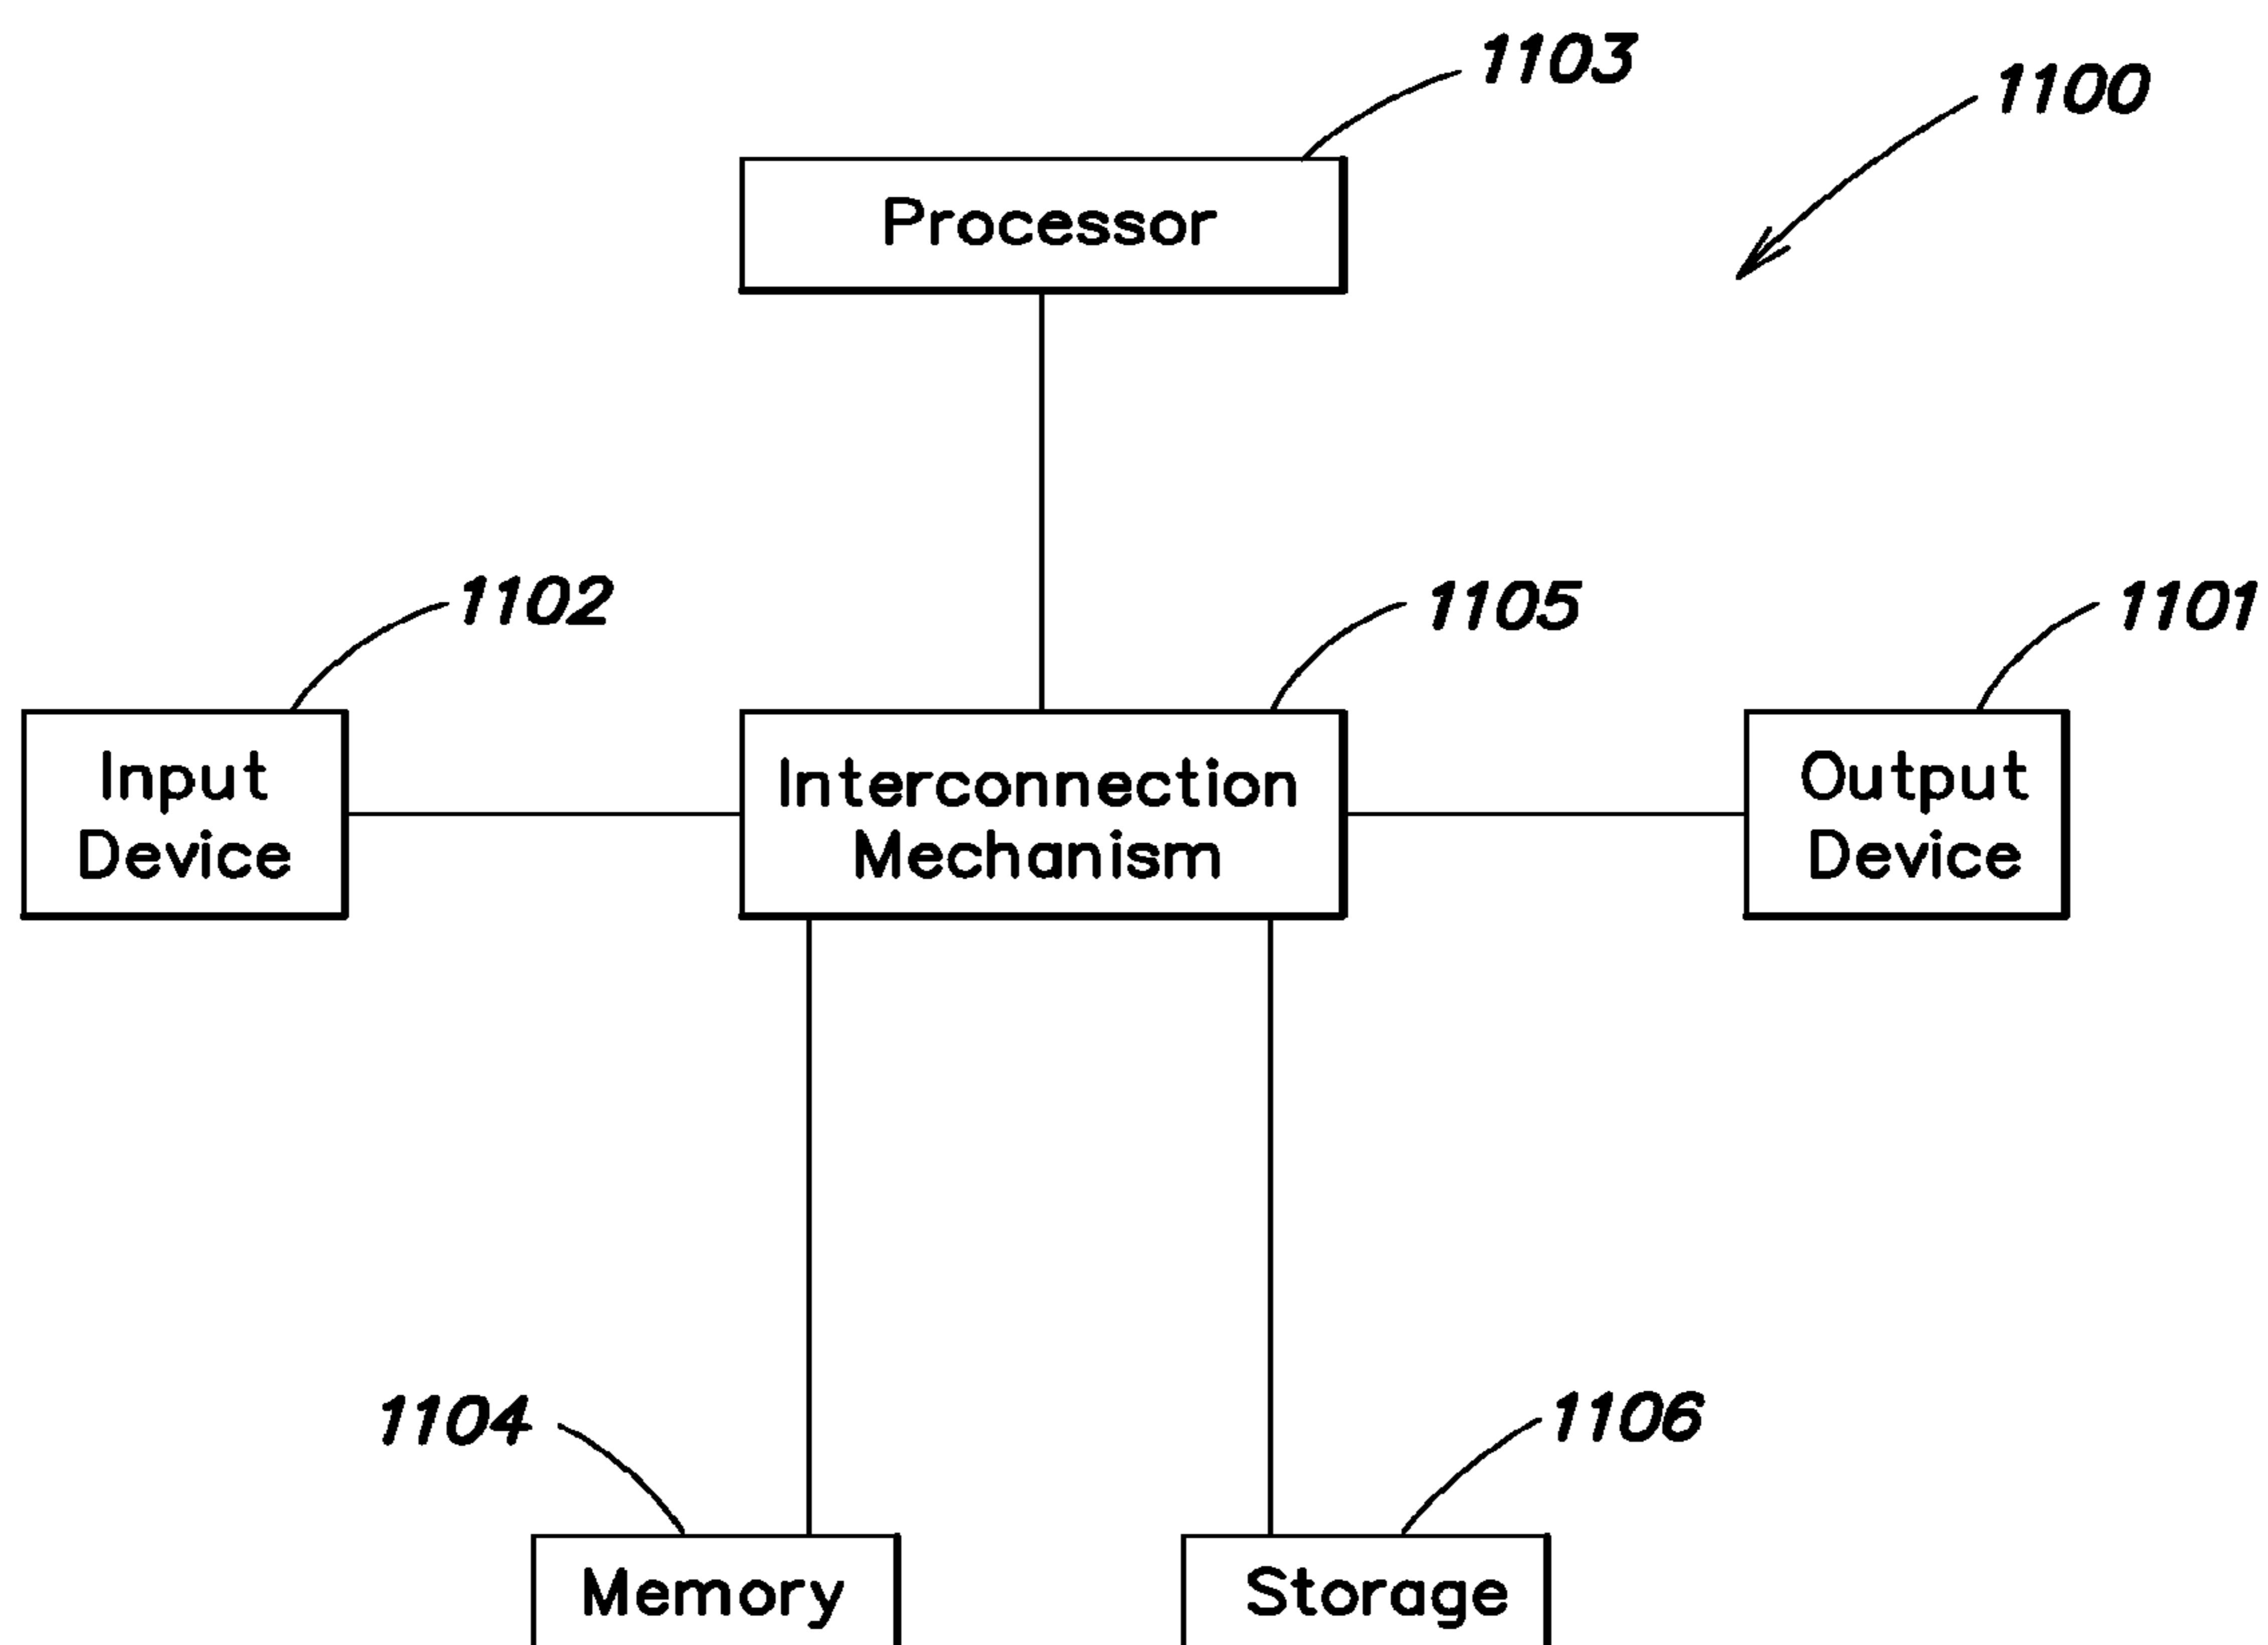


FIG. 11




FIG. 12




FIG. 13




FIG. 14




FIG. 15


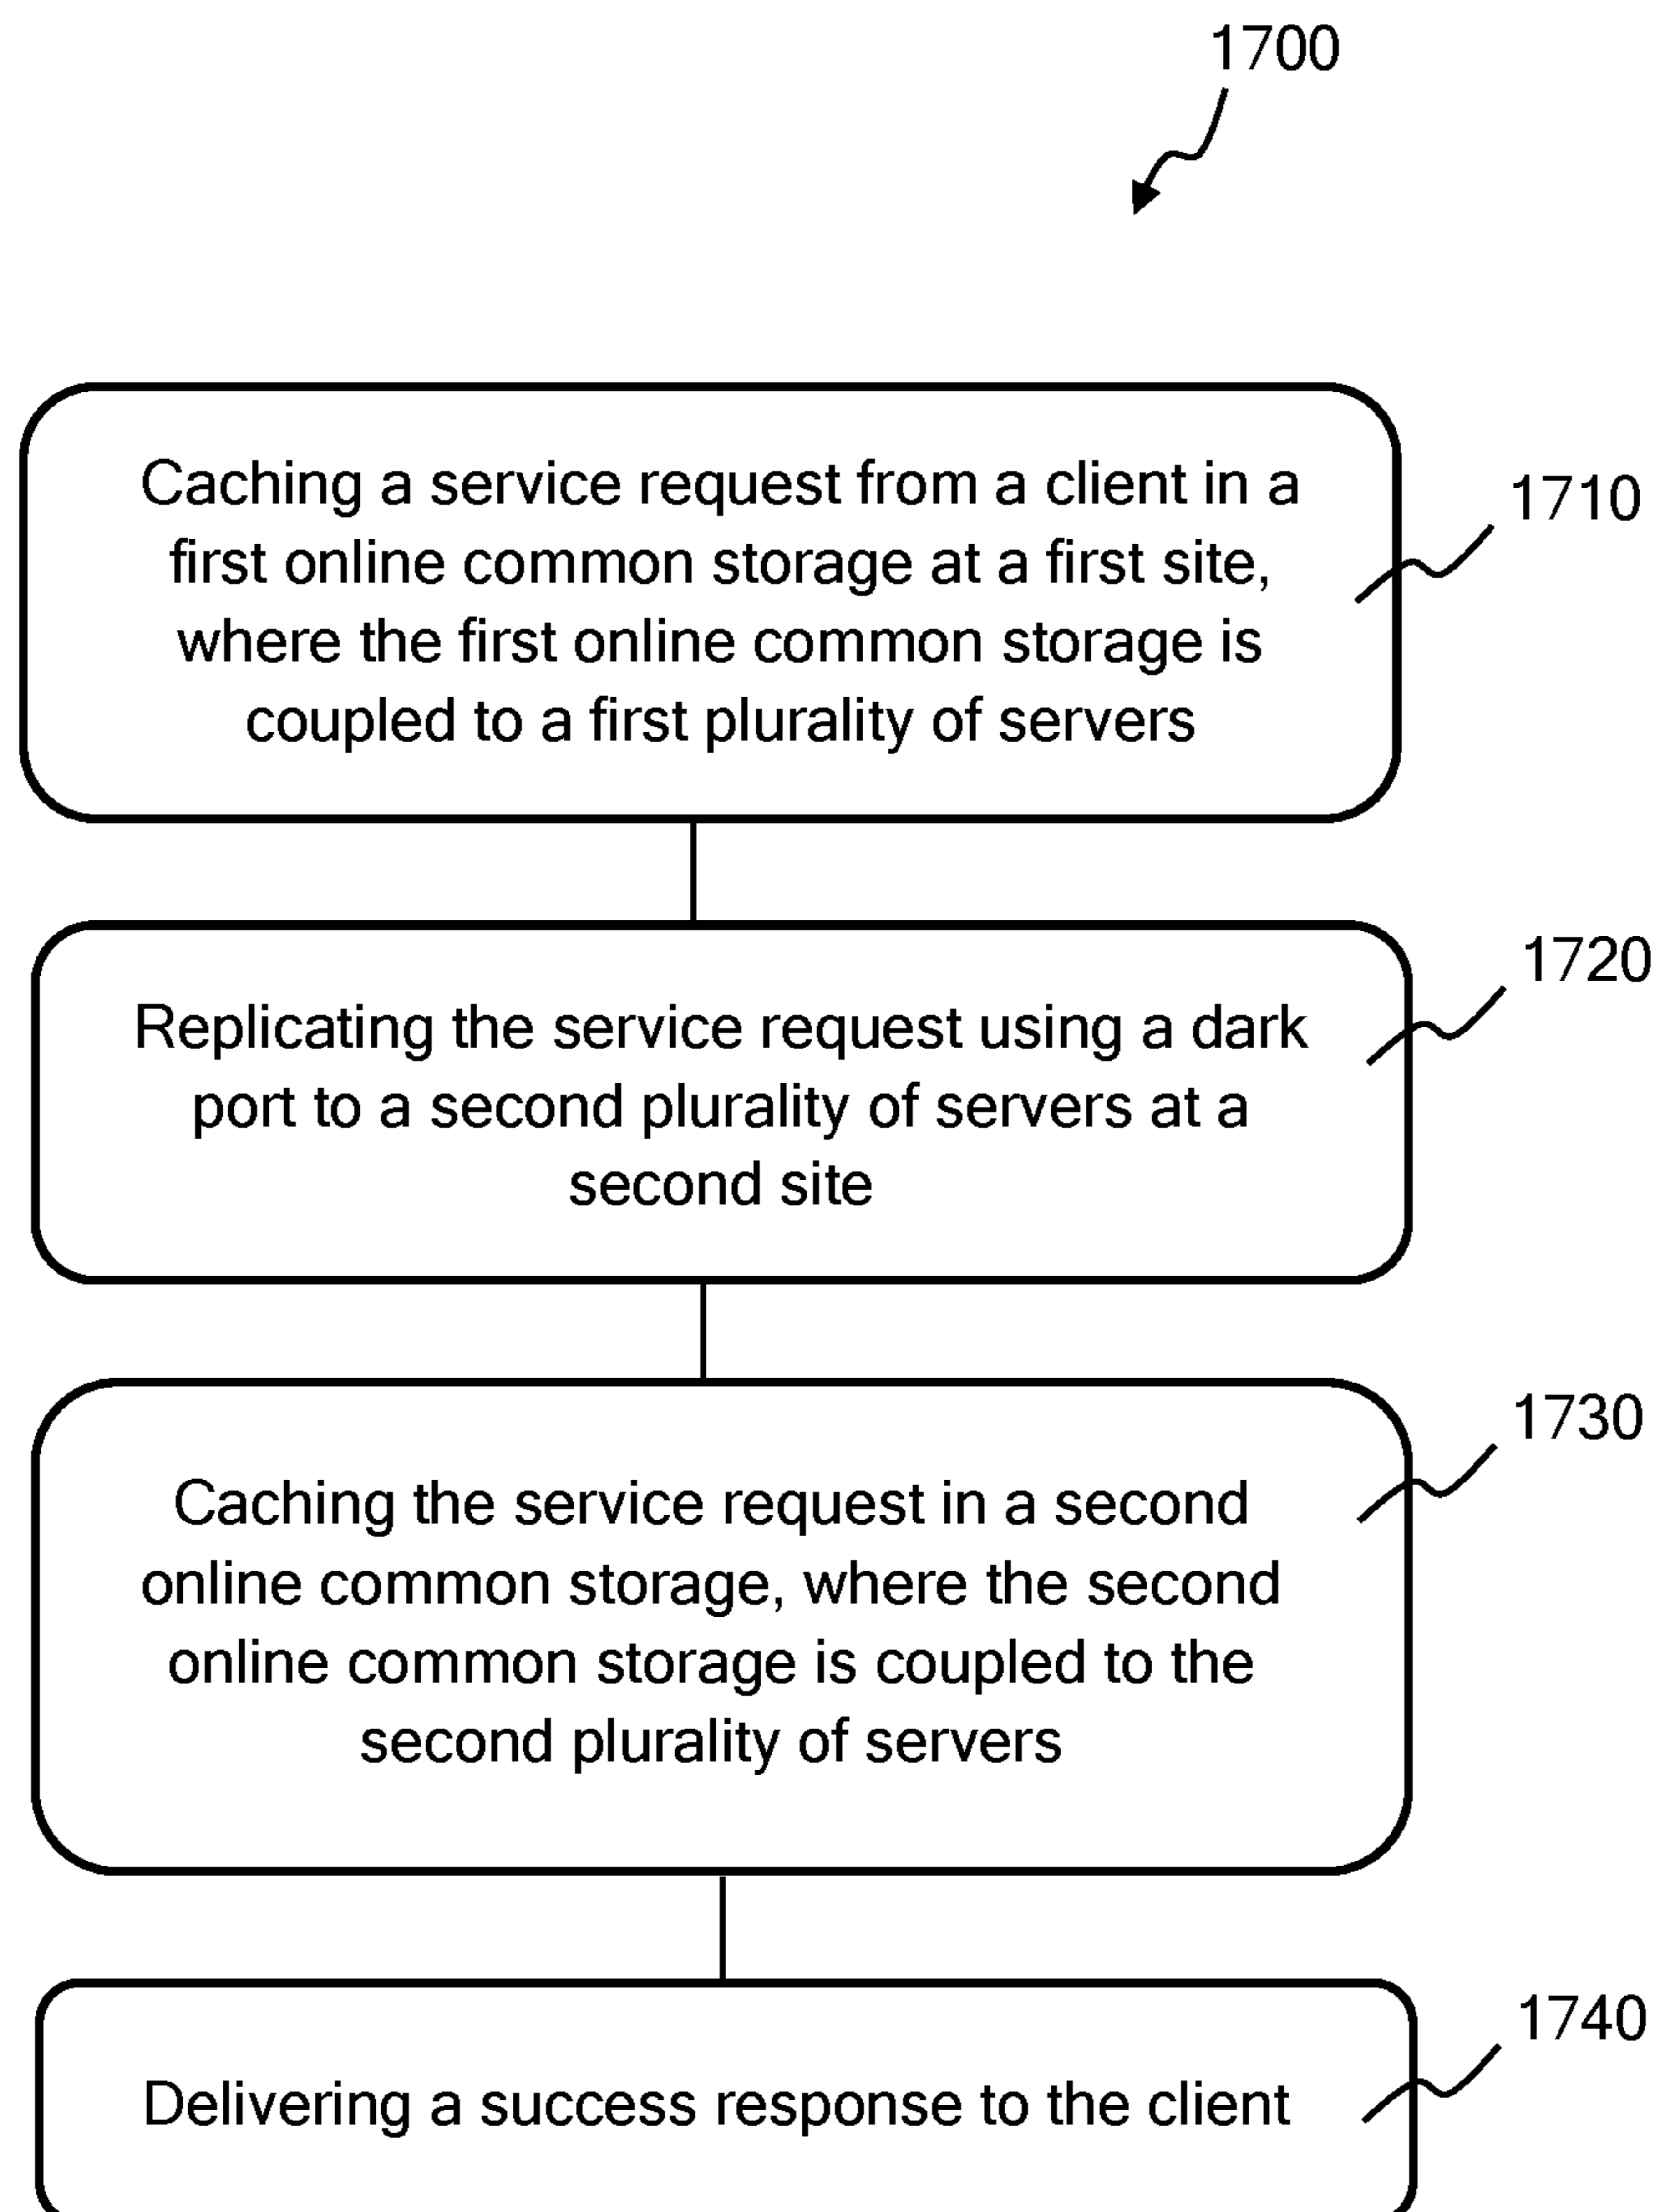
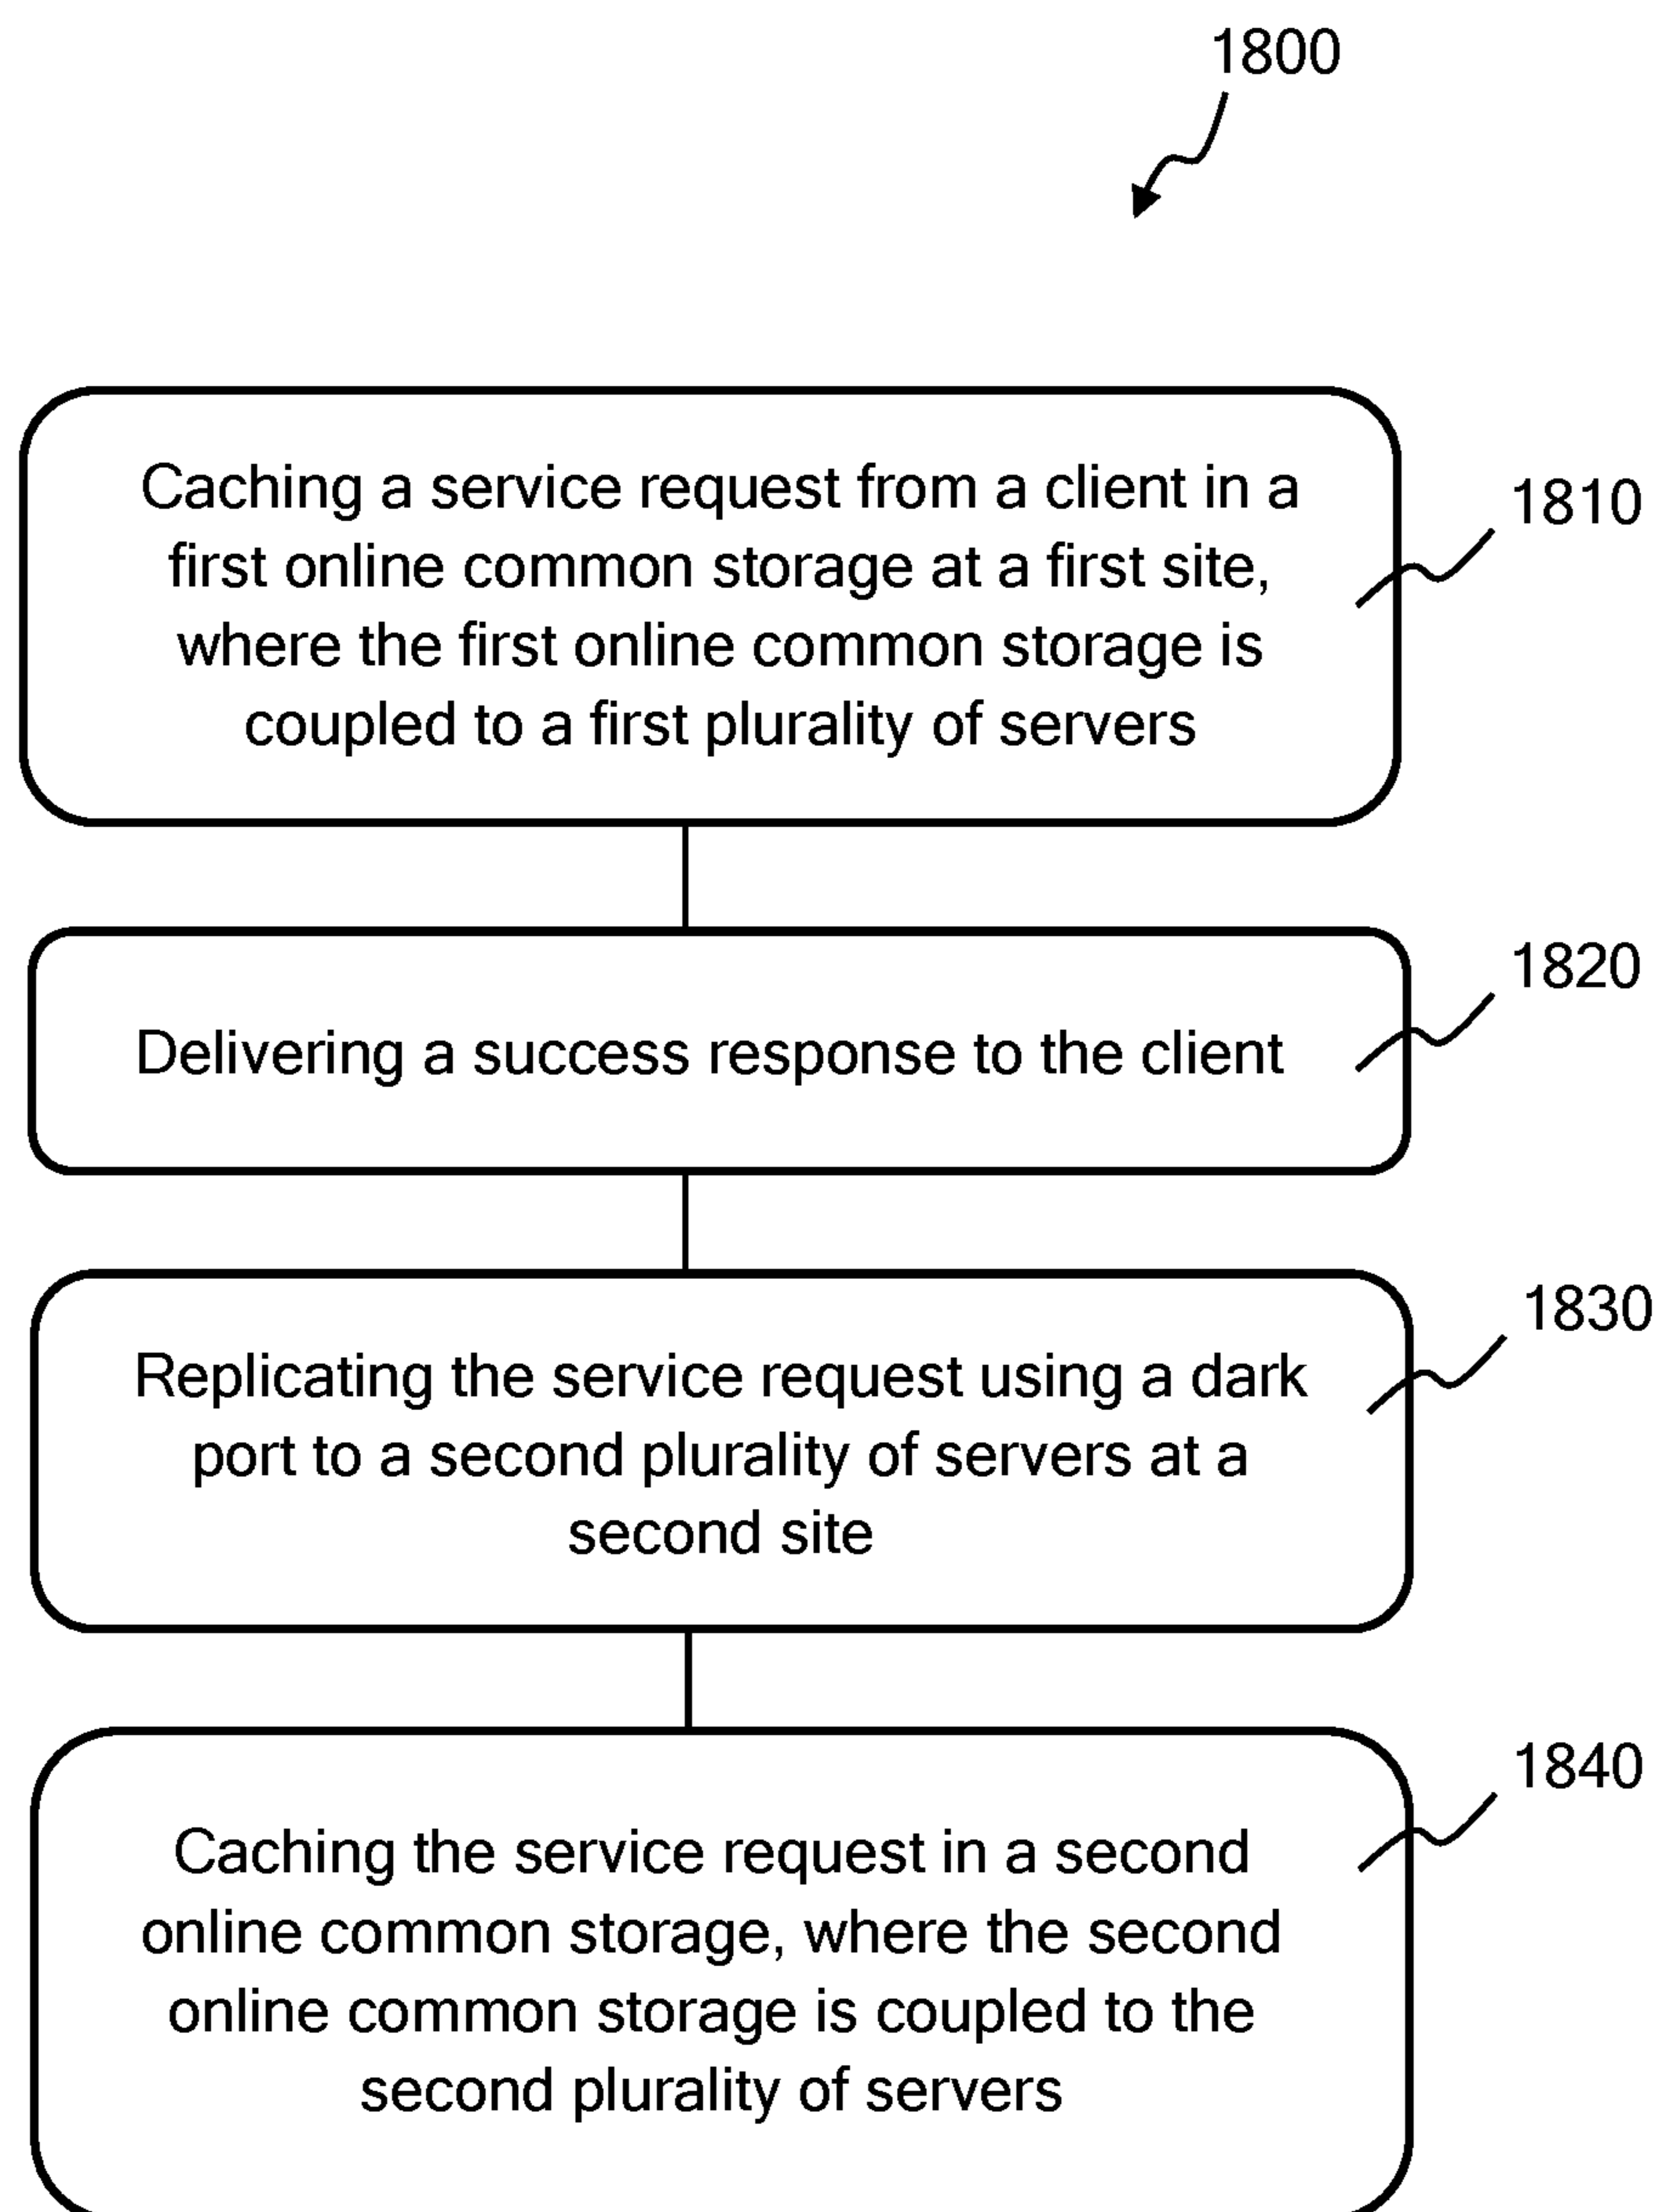
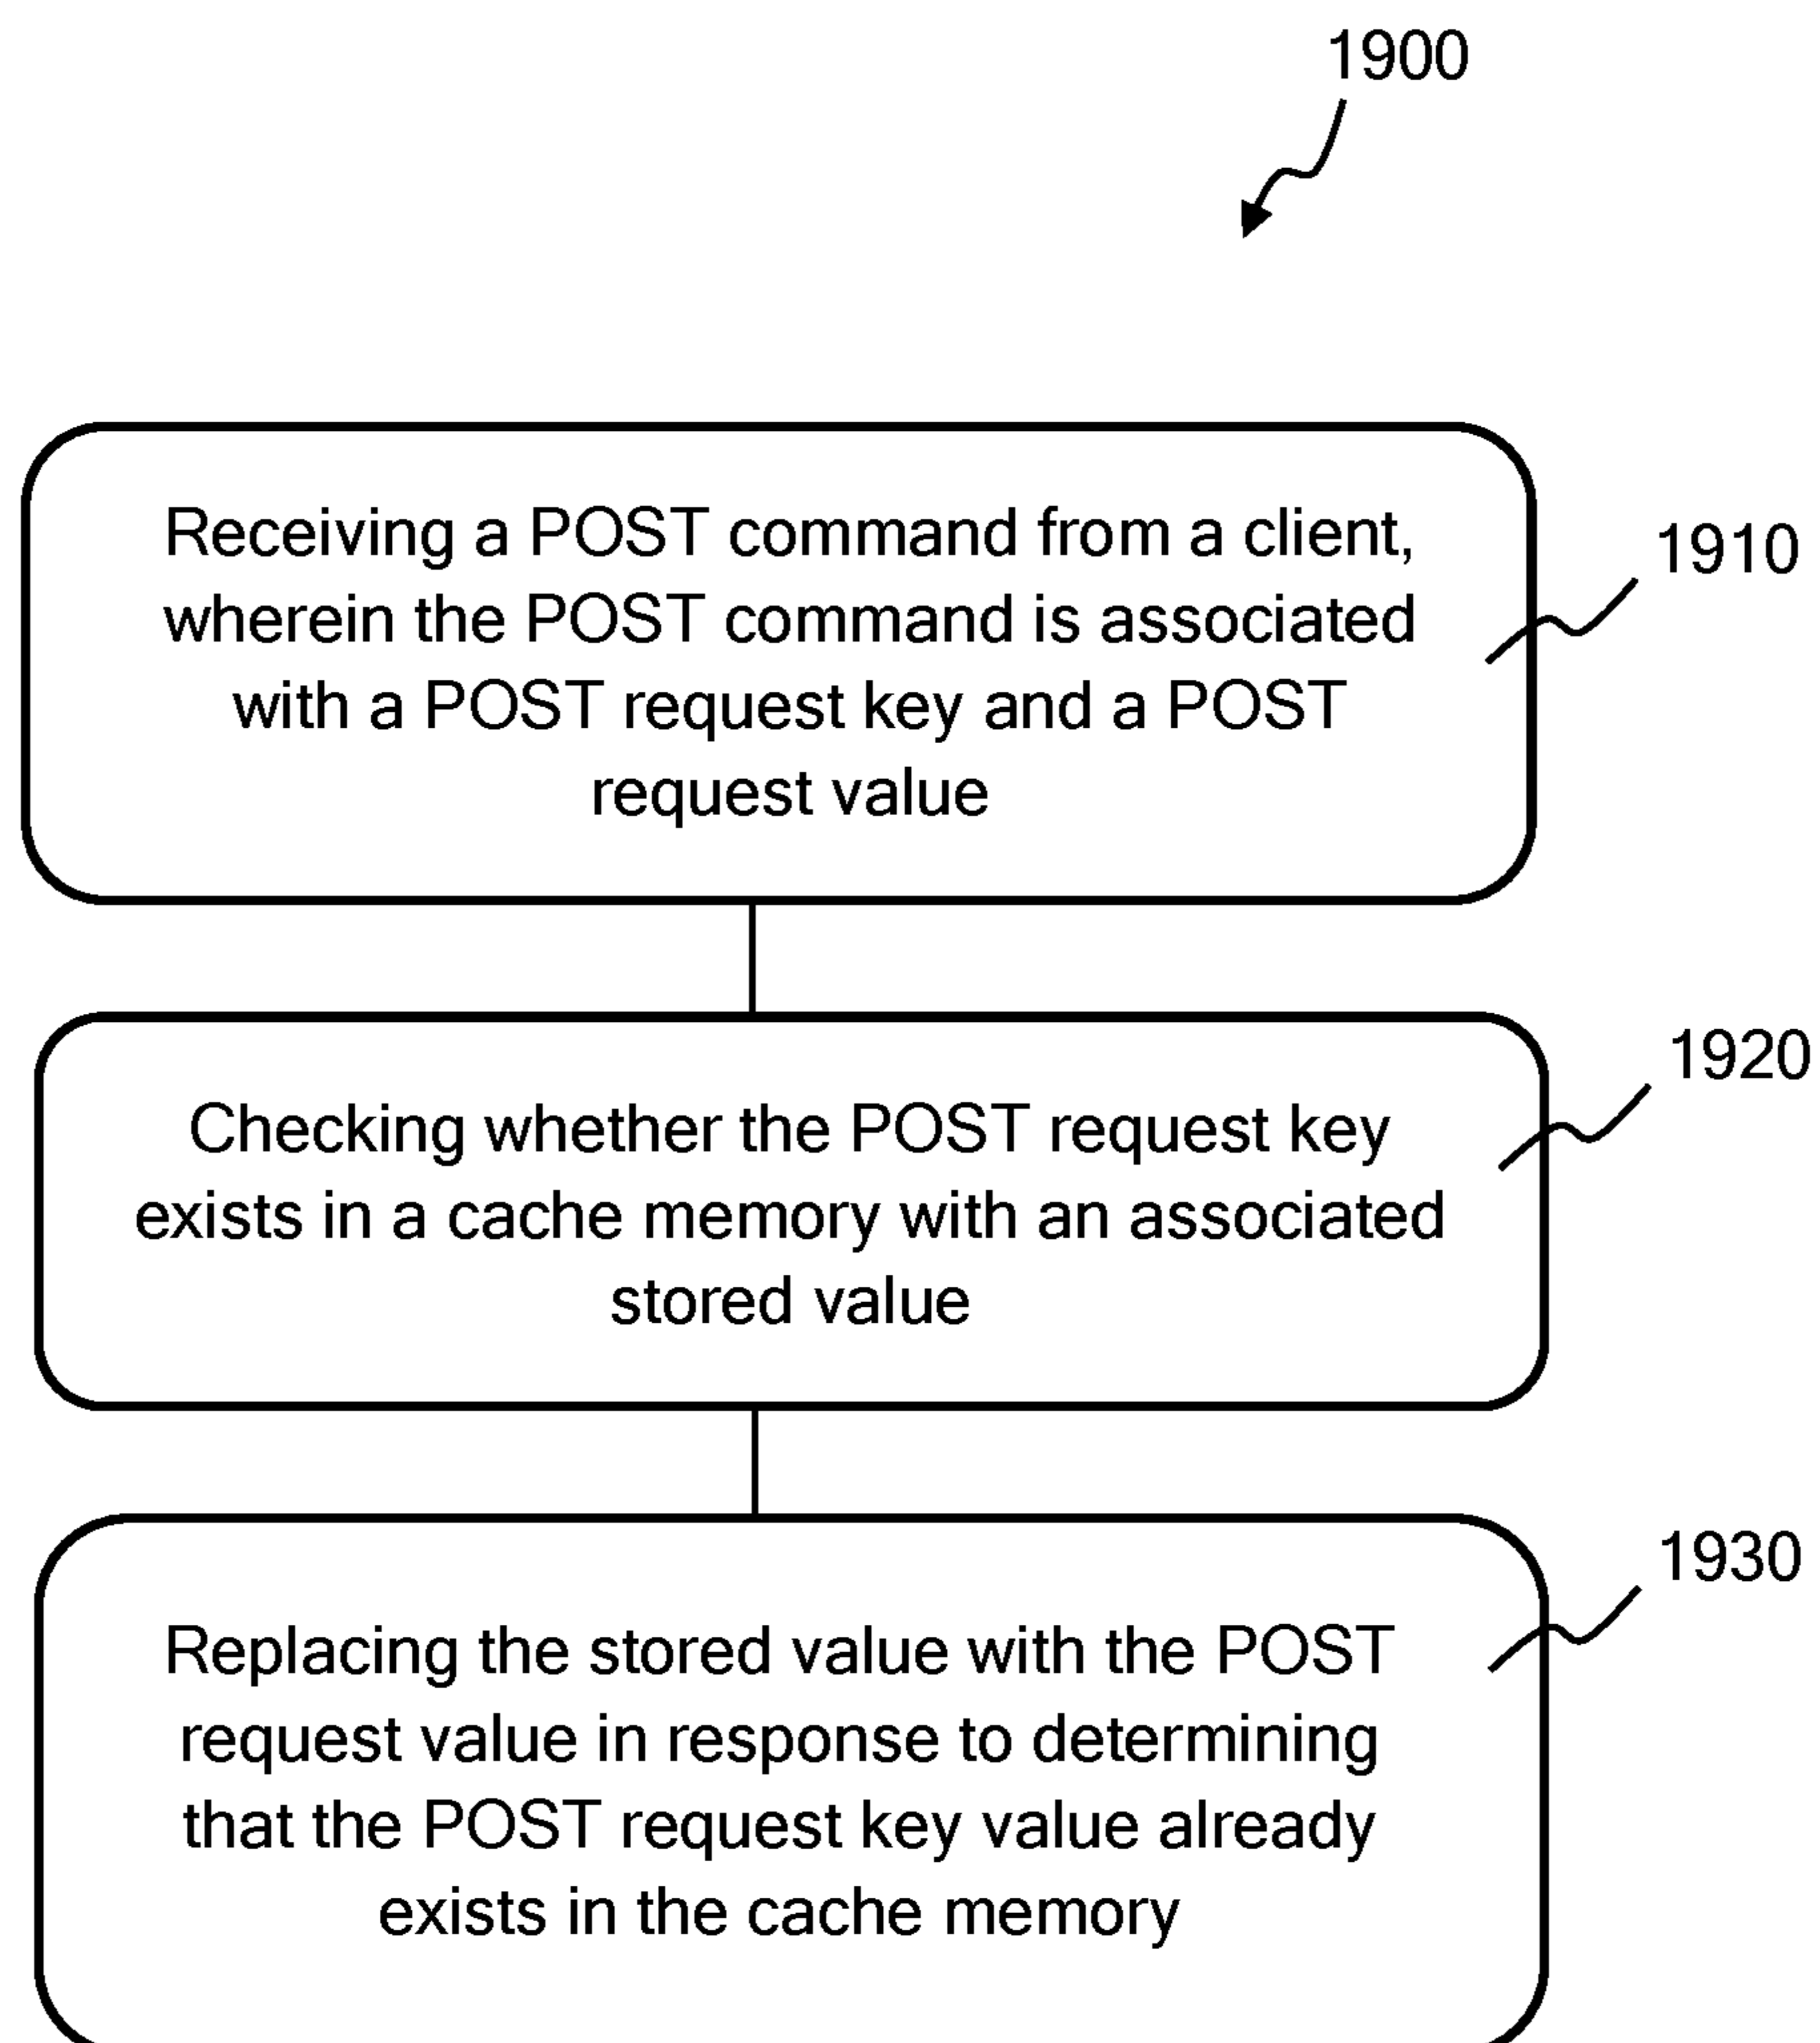


FIG. 16

**FIG. 17**

**FIG. 18**

**FIG. 19**

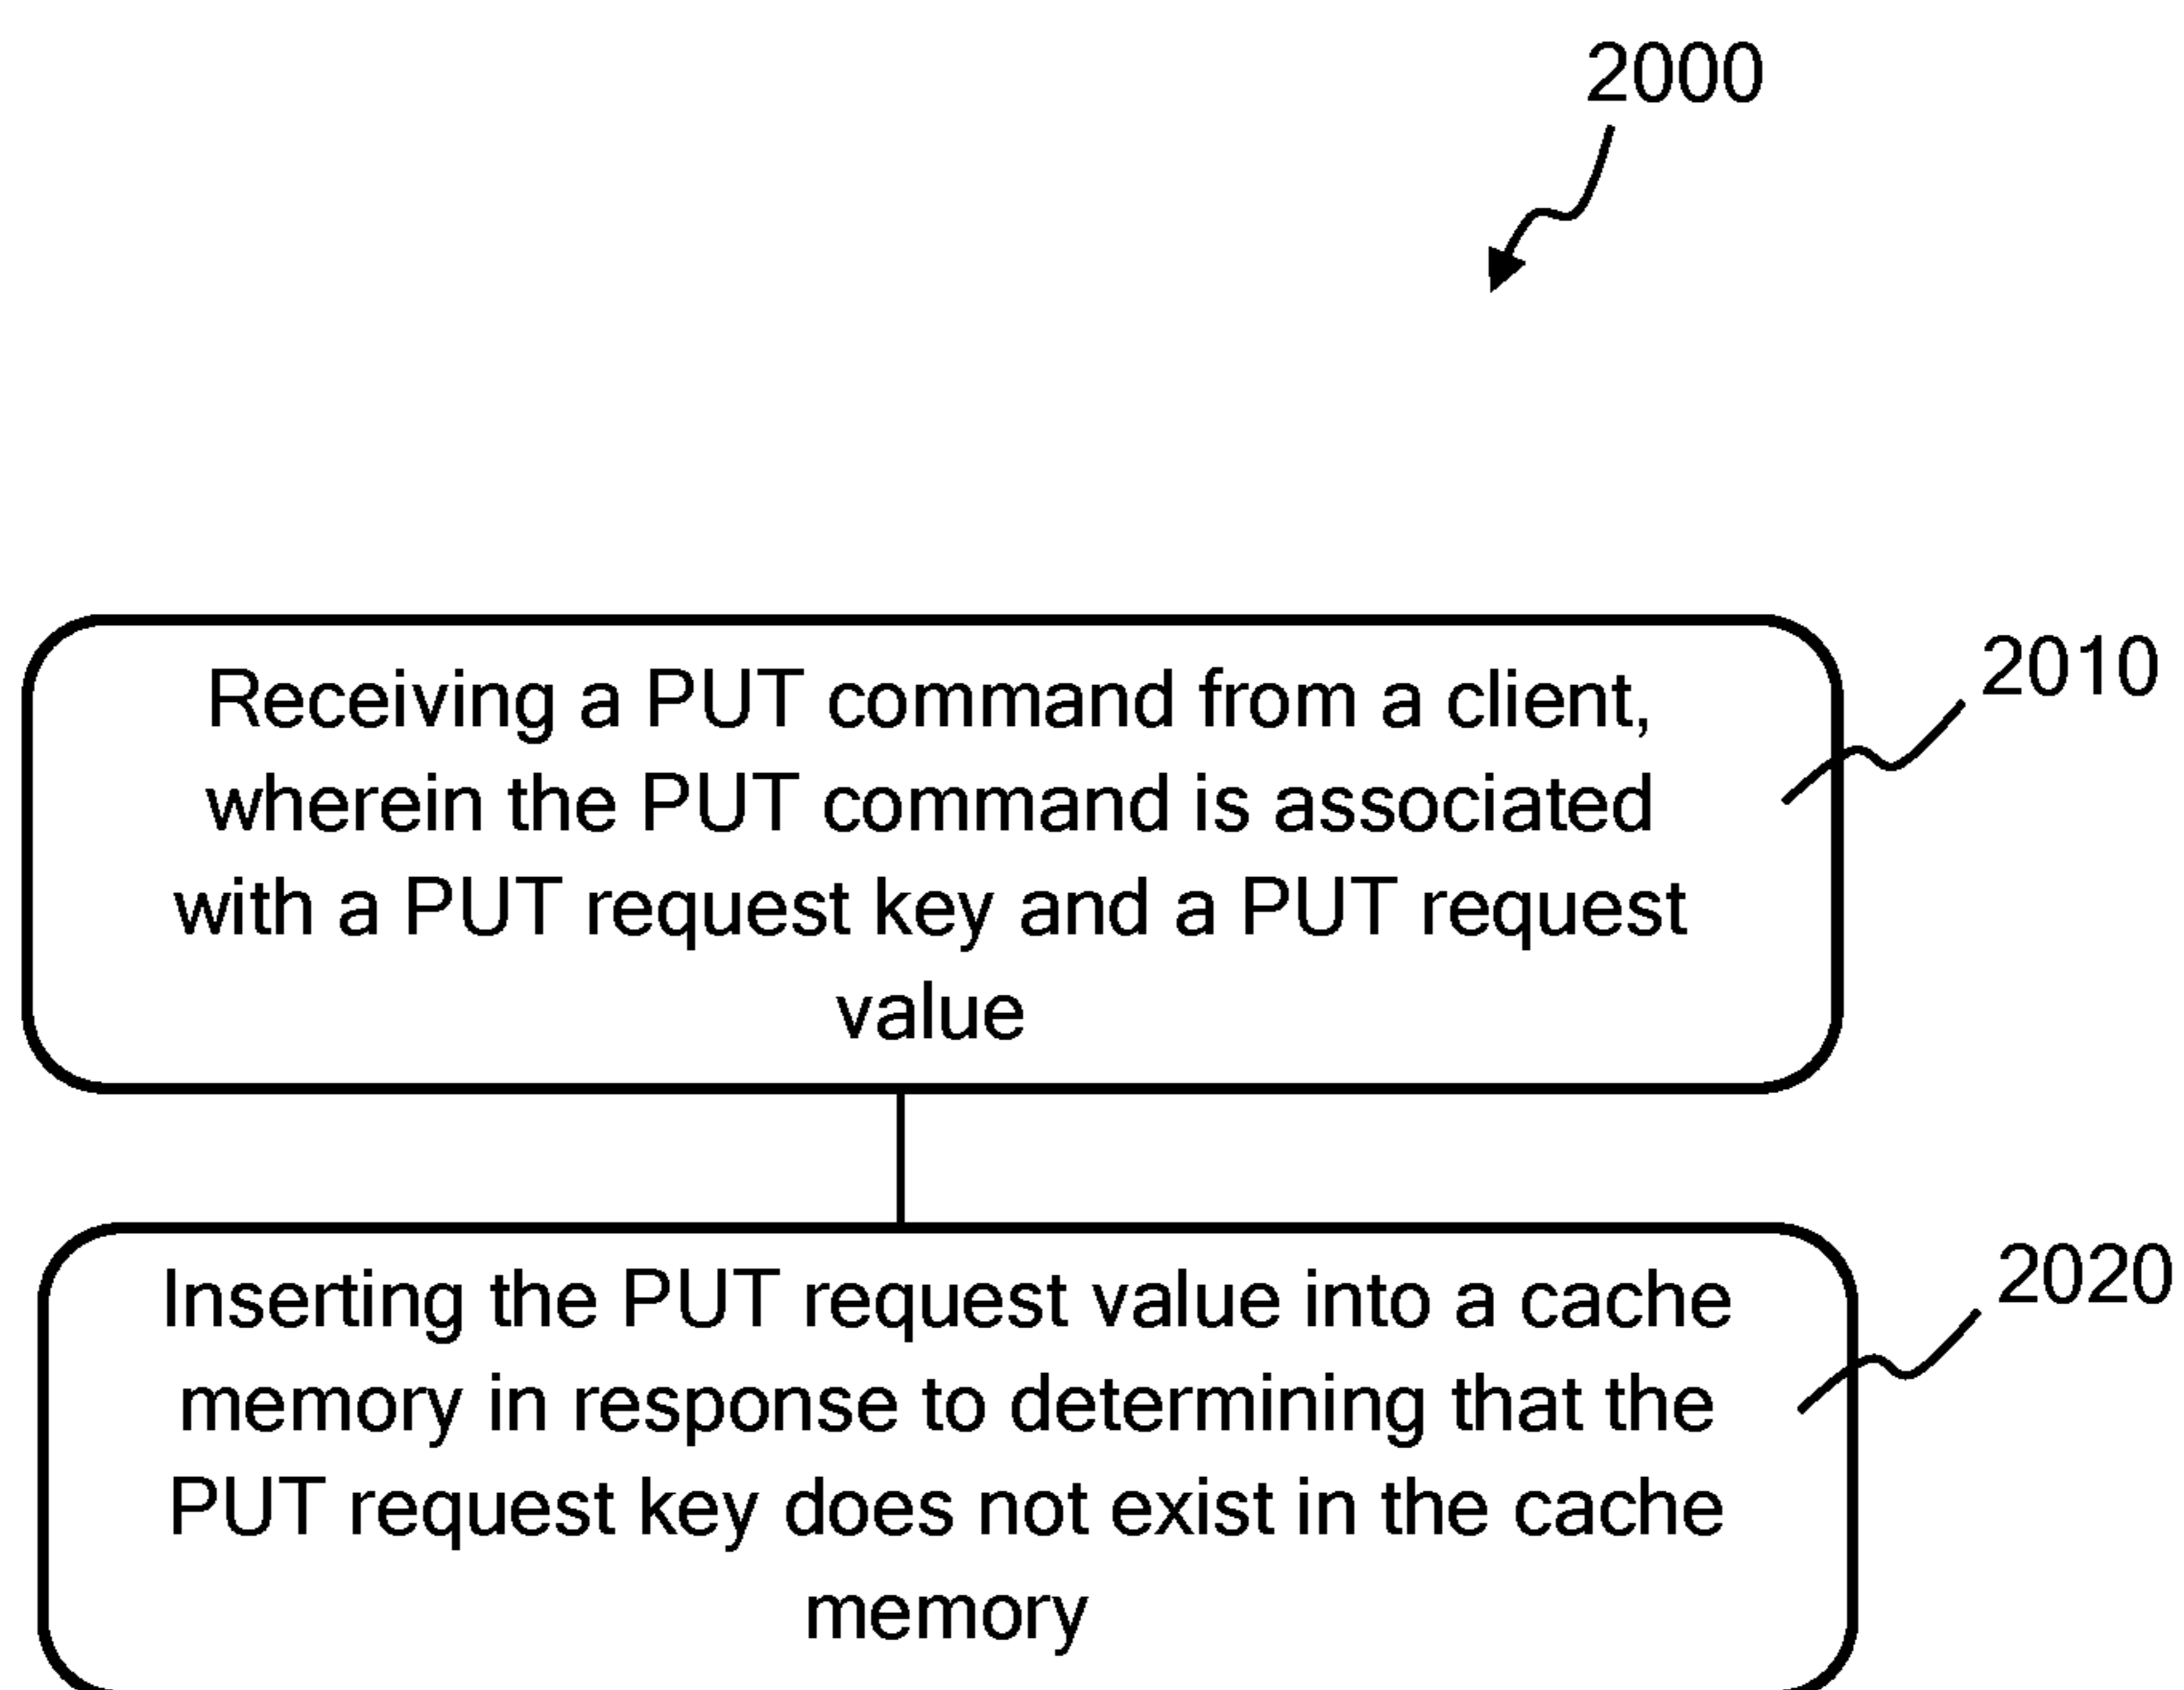
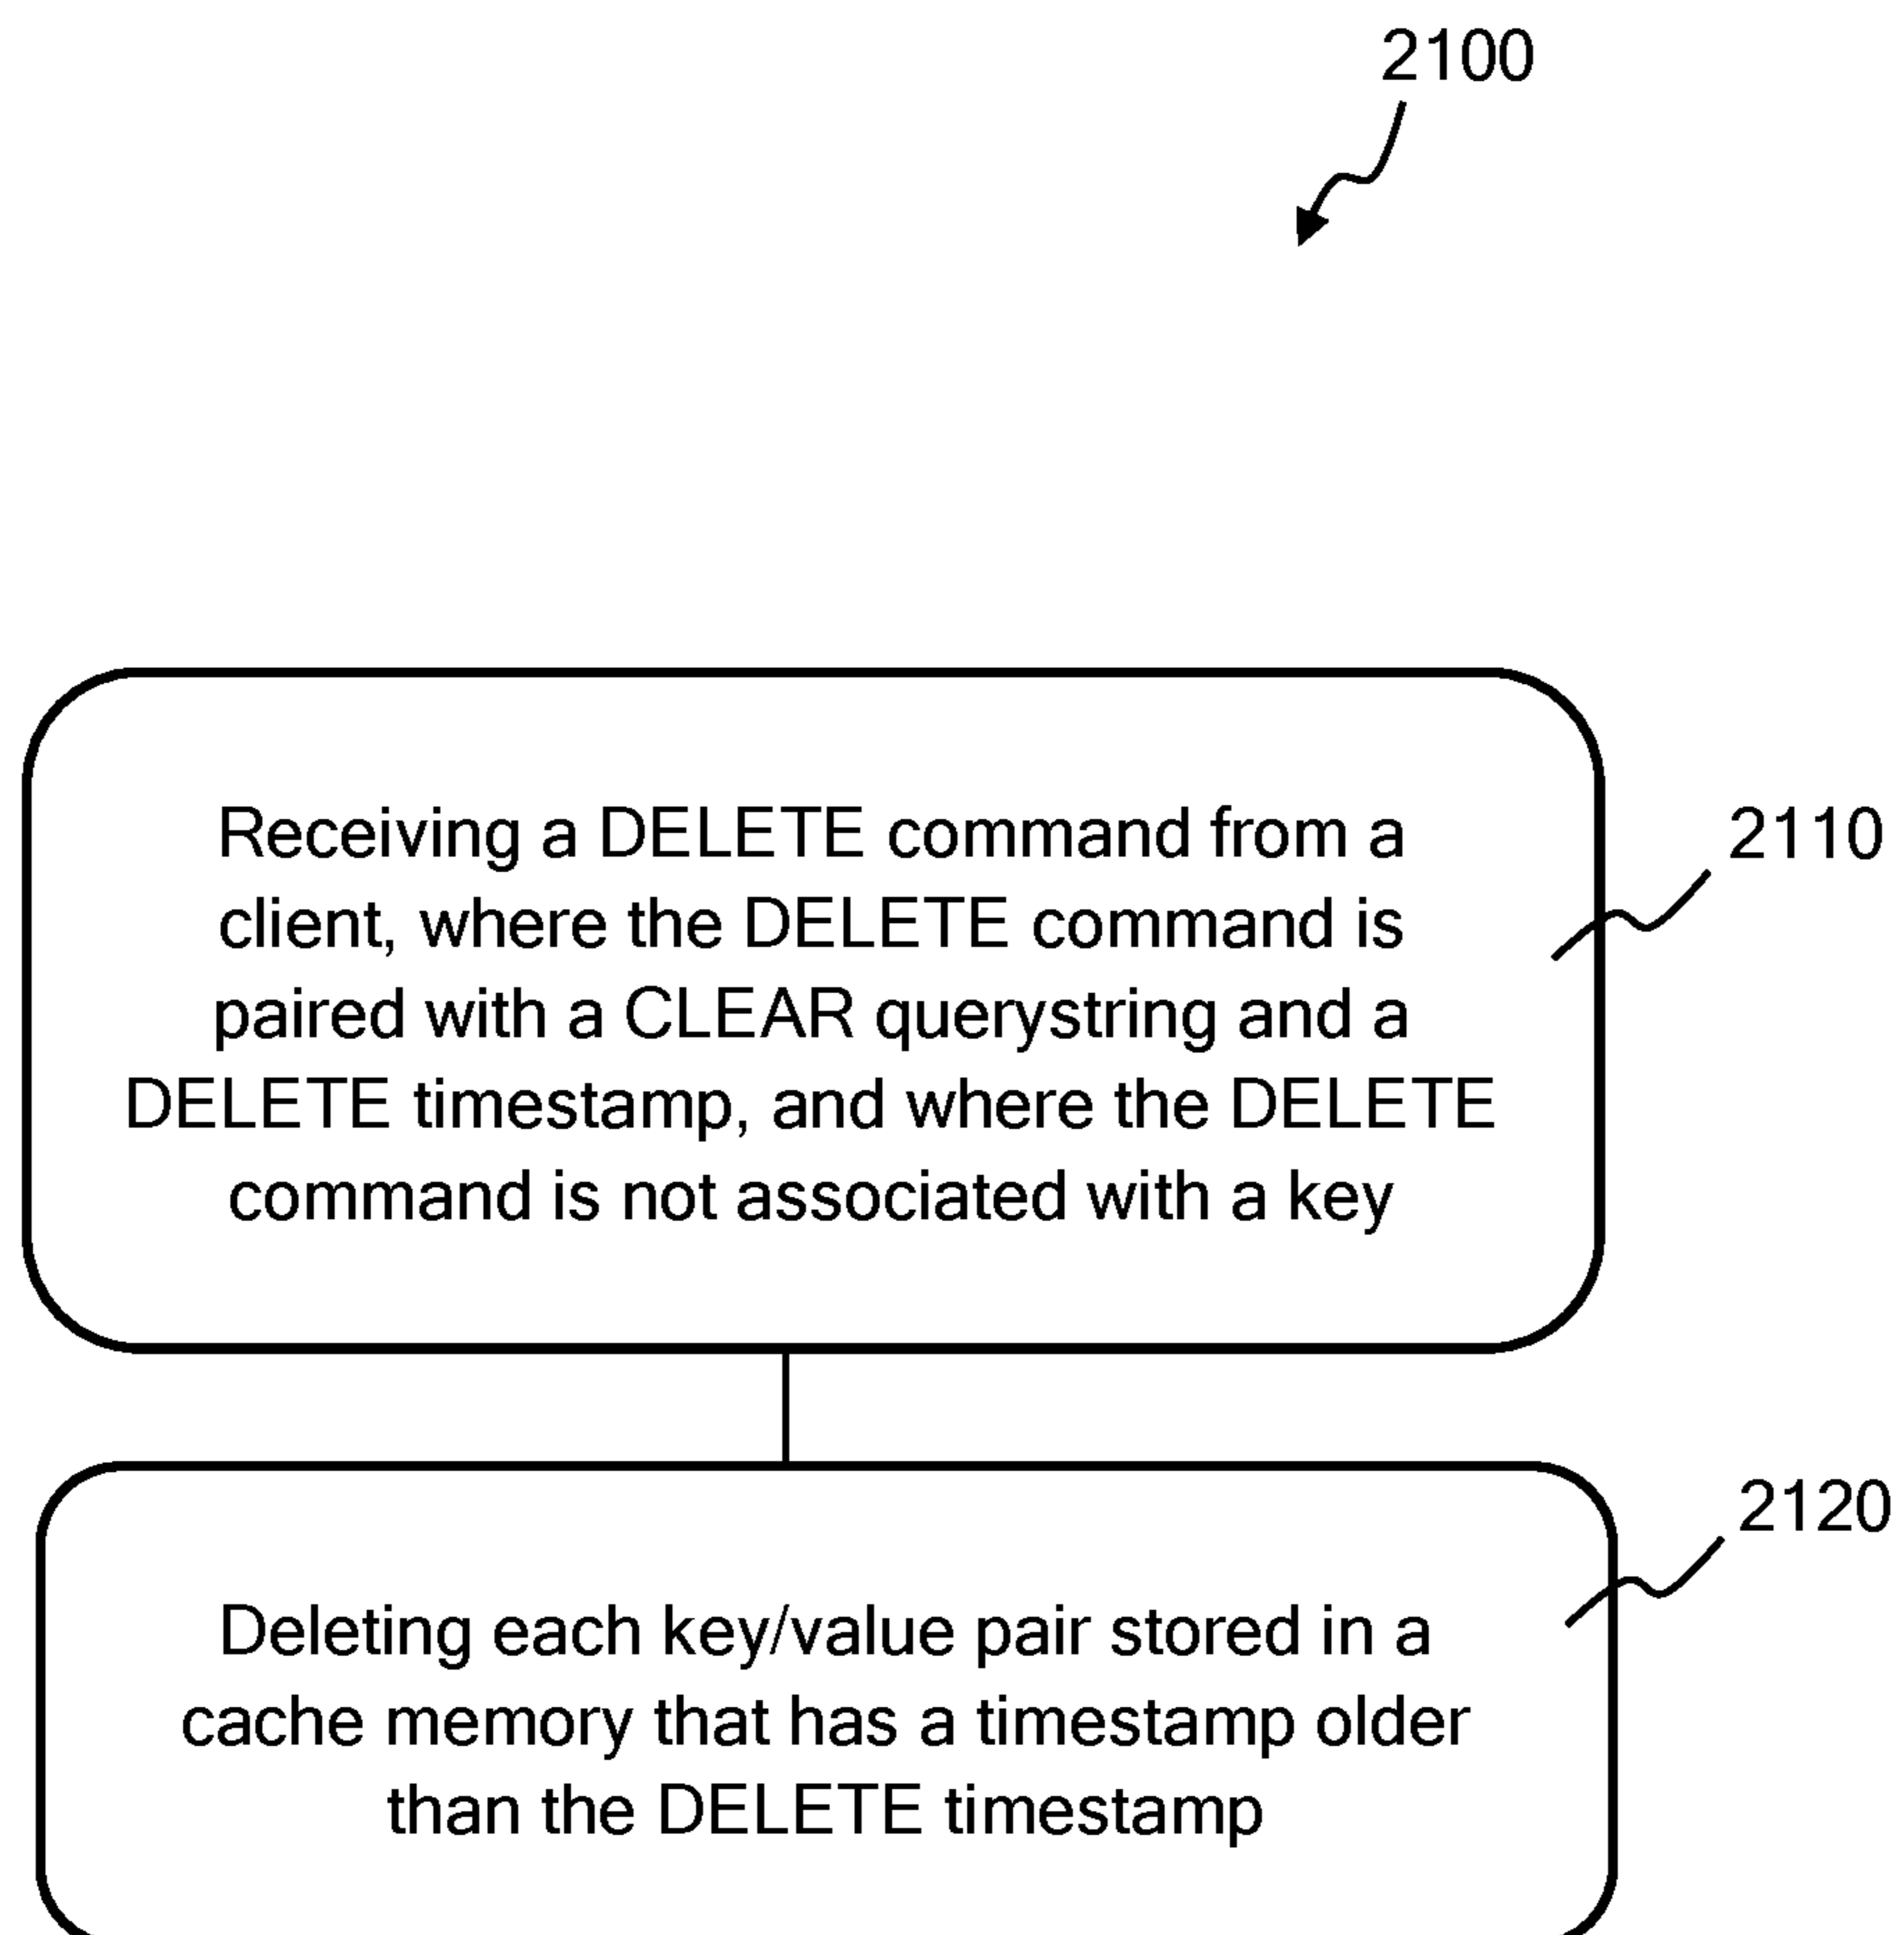


FIG. 20

**FIG. 21**




FIG. 22

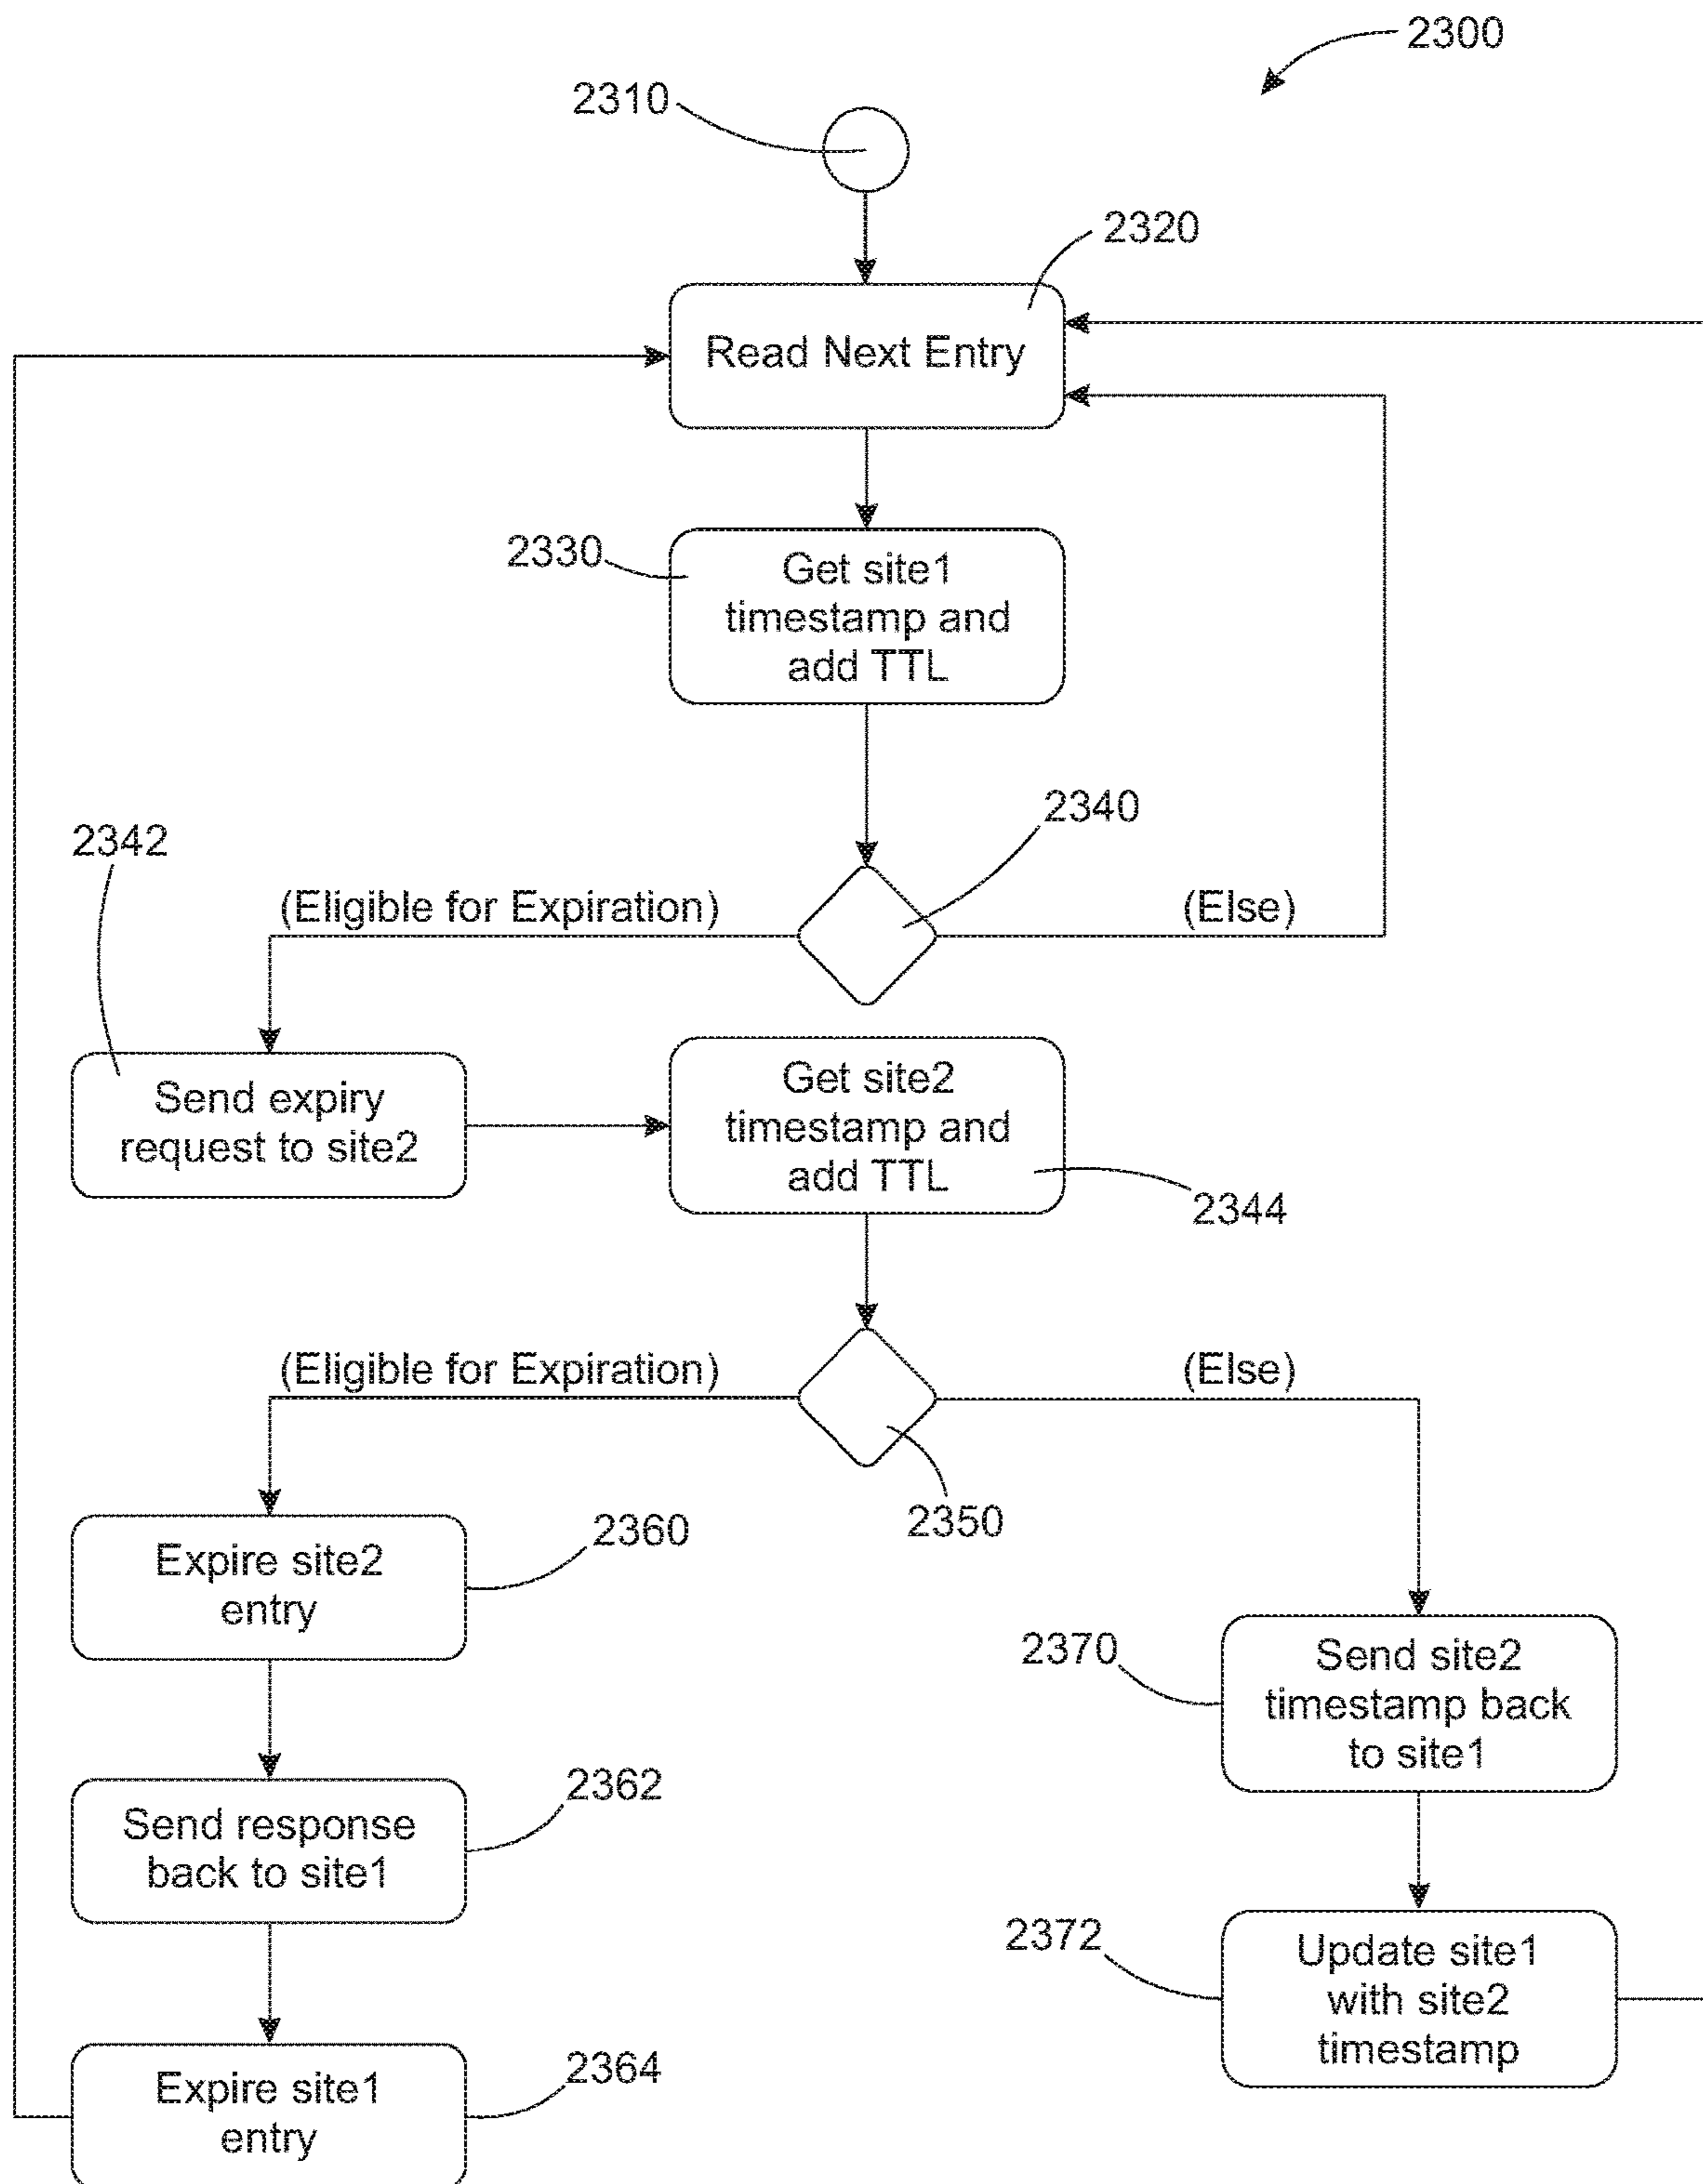
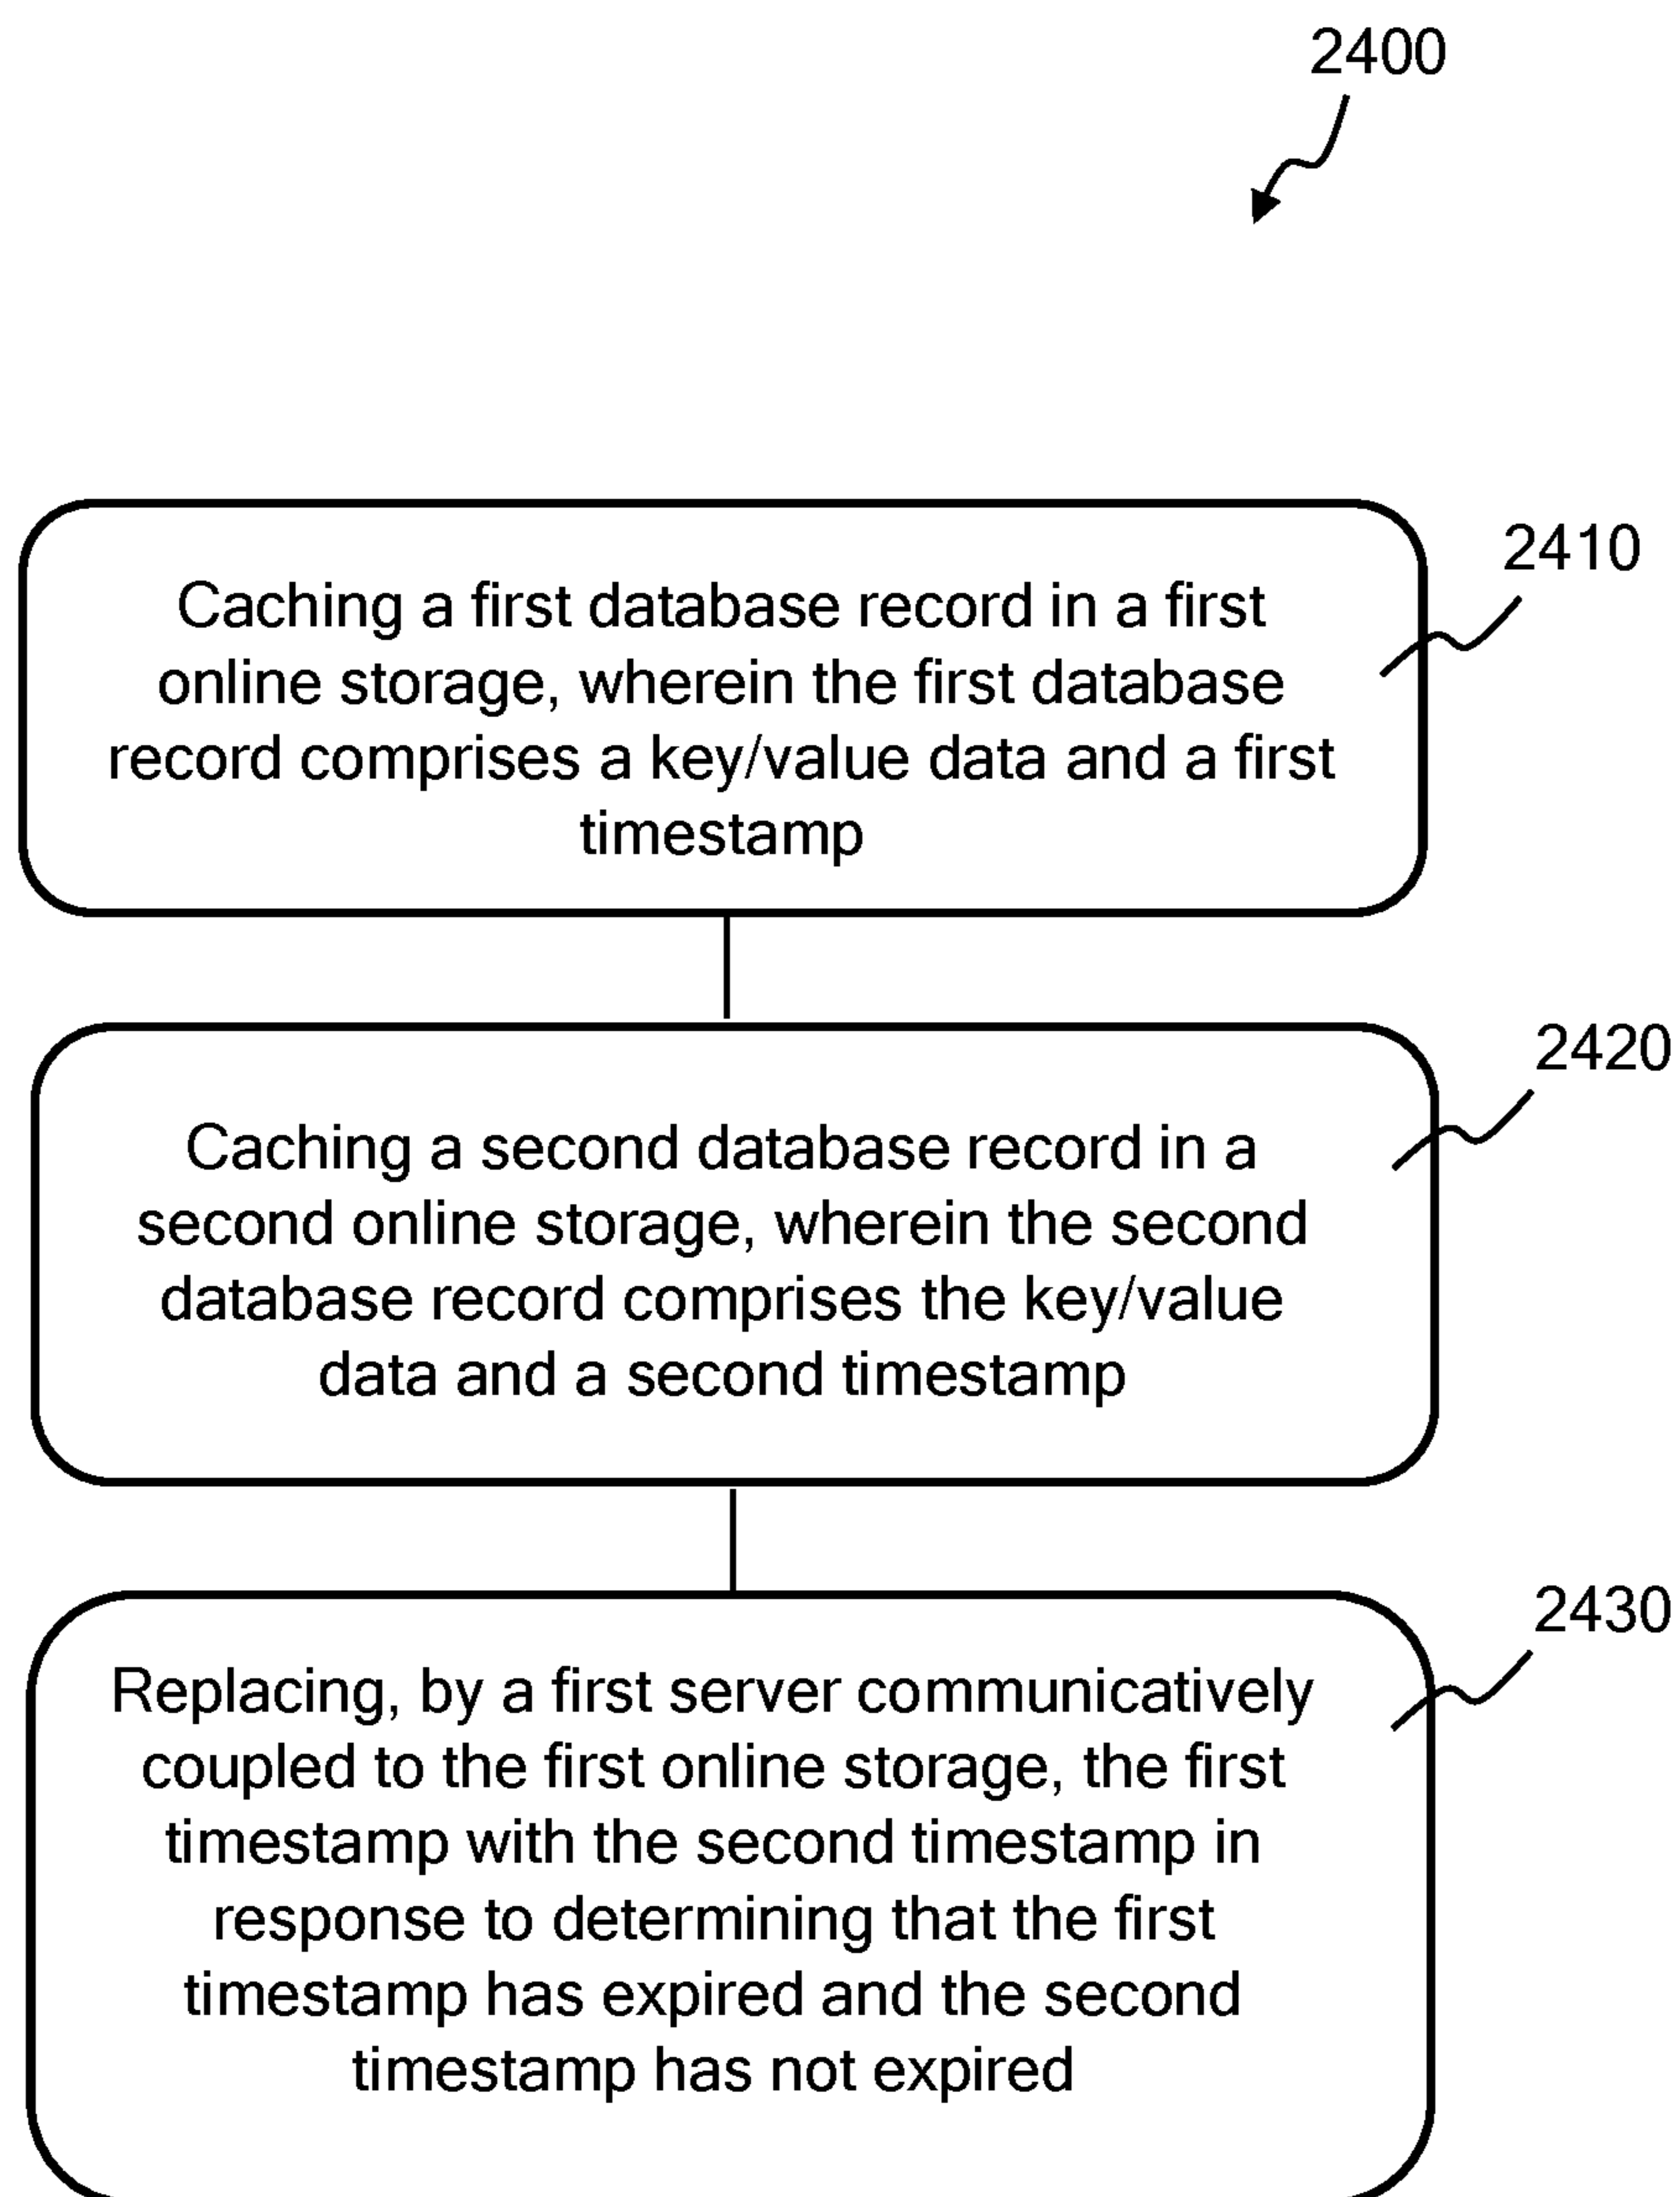


FIG. 23

**FIG. 24**

CACHING SYSTEM AND METHOD**CROSS REFERENCE TO RELATED APPLICATIONS**

This application is a continuation-in-part of U.S. patent application Ser. No. 15/143,953, filed May 2, 2016 and entitled “Caching System and Method”, which is hereby incorporated entirely herein by reference. This application is also a continuation-in-part of U.S. patent application Ser. No. 13/960,192, filed Aug. 6, 2013 and entitled “System and Method For Storing and Processing Web Service Requests”, which is hereby incorporated entirely herein by reference. U.S. patent application Ser. No. 15/143,953 is a continuation-in-part of U.S. patent application Ser. No. 13/960,266, filed Aug. 6, 2013 and entitled “System and Method For Processing Web Service Transactions using Timestamp Data”, and a continuation-in part of U.S. patent application Ser. No. 13/960,192, filed Aug. 6, 2013 and entitled “System and Method For Storing and Processing Web Service Requests” which are hereby incorporated entirely herein by reference.

BACKGROUND OF THE INVENTION**Technical Field**

This invention relates to caching systems, and more specifically to a system and method for processing web-based service requests.

State of the Art

There are many different caching system and methods for processing web-based service requests. For instance, there are systems for load sharing service requests among a number of server-based systems. Conventionally, systems that fail during the processing of web-based requests generally do not fail gracefully, and service errors occur as a result. Accordingly, it would be beneficial to have a web-based caching service that is capable of being scaled depending on particular clients or applications.

DISCLOSURE OF THE INVENTION

According to one aspect, it is appreciated that it may be useful and particularly advantageous to provide a system that is capable of servicing one or more web-based service requests. According to one embodiment, a platform is provided that allows web-based service requests to be stored and served in a more reliable manner. For instance, a caching service may be provided that is capable of storing and tracking received requests, and, responsive to failures in a distributed computer system, is capable of transitioning those service requests to other distributed computer system resources. Conventionally, systems that fail during the processing of web-based requests generally do not fail gracefully, and service errors occur as a result.

Further, it would be beneficial to have a web-based caching service that is capable of being scaled depending on particular clients or applications. For instance, a cloud-based platform may be provided that hosts applications that support multiple users, and each service for a particular provider may be tailored to meet the needs of the application. According to various embodiments, the caching service that is provided to support such applications may be adjusted to support certain fault tolerance levels, response times,

failover scenarios, and data replication and backup requirements. Optimally, the caching service exhibits high-performance, is highly available, efficient, reliable, scalable and cost-effective.

5 According to one embodiment, the caching service may be configured in one of a number of different environments. For instance, the caching service may be configured as an active/single environment where a single caching service entity responds to one or more web service requests. In another environment, a caching service may be configured as an active/active environment where two or more caching service entities receive and respond to one or more web service requests. In yet another implementation, a caching service may be configured as an active/standby system, where two or more caching service entities receive one of the web service requests, but a single entity is responsible for committing transactions.

Such a system will be beneficial when supporting one or more data centers that include systems that serve as backup systems that respond to service web-service requests. For instance, active/active configurations may be used to cache across multiple data centers for the purpose of performing synchronous replication. In another example, active/standby configurations may be used to cache requests across multiple data centers for the purpose of performing asynchronous replication to a “hot” standby data center.

According to another embodiment, the caching system will support representational state transfer (REST) and simple object access protocol (SOAP) service calls. According to another embodiment of the present invention, a caching system supports hypertext transfer protocol (HTTP) and hypertext transfer protocol secure (HTTPS) service calls. In another embodiment, the service calls may be encrypted (e.g., via secure socket layer/transport layer security (SSL/TLS) encryption). In one embodiment, the service requests may include commands such as POST, GET, PUT and DELETE requests. In one implementation, basic security may be provided for REST service requests (e.g., by providing support for userid/password authentication and authentication using resource access control facility (RACF)). In one implementation, basic security may be provided for SOAP service requests (e.g., by providing support for userid/password authentication using SOAP header, and authentication using RACF).

5 In yet another embodiment, a distributed method for updating timestamp information related to stored data is provided. For instance, it is appreciated that timestamp information for data may need to be made consistent across multiple datasets (e.g., located at different datacenters). It is realized that in high-volume transaction systems, it may not be feasible to adequately transmit timestamp information in an effective manner between systems. For instance, excessive network traffic would be created by synchronizing such timestamp information created by multiple updates and data accesses. Therefore, it would be preferable to permit timestamp updates that minimize network traffic. Further, a capability may be provided that permits timestamp information to be maintained for a data element based on when the data was last accessed. For instance, most data elements only include information that identifies when a data element was last updated, not when it was last accessed.

According to one aspect, a system for processing web service requests is provided. The system comprises a server configured to receive and process web service requests, the server comprising a plurality of virtual computer systems adapted to service received web service requests; a logical storage system coupled to the plurality of virtual computer

systems, wherein each of the plurality of virtual computer systems shares a common storage that is adapted to store the received web service requests; and a request handler element adapted to distribute a web service request to at least one of the plurality of virtual computer systems for processing. In one embodiment, the web service requests include at least one of a group comprising a REST service request; and a SOAP service request. In one embodiment, the server configured to receive and process web service requests is operable in at least one of a group of operating modes, the group comprising an active/single mode; an active/standby mode; and an active/active mode.

In one embodiment, the system further comprises a load balancing element adapted to distribute the web service requests among a plurality of server systems. In one embodiment, the plurality of virtual computer systems is located within a partition. In one embodiment, a web service application is assigned to a particular partition. In one embodiment, the system further comprises transaction server components that are adapted to process web service transactions. In one embodiment, the logical storage system further comprises a common database shared by the plurality of virtual servers upon which web service requests are transacted. In one embodiment, the system further comprises an entity that monitors an expiration of a database record associated with at least one web service request.

According to one aspect, the system further comprises corresponding entities that execute among at least two of the plurality of virtual computer systems, the entities being adapted to compare timestamps associated with the database record associated with the at least one web service request. In one embodiment, the entity is adapted to delete the database record associated with the at least one web service request. In one embodiment, the at least two of the plurality of virtual computer systems execute within separate computer systems. In one embodiment, at least two of the virtual computer systems are located in different data centers.

According to one aspect, a method for processing web service requests is provided. The method comprises receiving, by a load sharing entity, a plurality of web service requests from one or more client systems; storing, in a common storage location, the received plurality of web service requests; assigning at least one virtual computer system to process at least one of the plurality of web service requests stored in the common storage location; and providing a response to the one or more client systems that generated the at least one of the plurality of web service requests. In one embodiment, the at least one of the plurality of web service requests includes at least one of a group comprising a REST service request; and a SOAP service request. In one embodiment, the method further comprises an act of operating the at least one virtual computer system in at least one of a group of operating modes, the group comprising an active/single mode; an active/standby mode; and an active/active mode. In one embodiment, the method further comprises an act of distributing the plurality of web service requests among a group of virtual computer systems. In one embodiment, the group of virtual computer systems is located within a partition. In one embodiment, the method further comprises an act of assigning a web service application to the partition. In one embodiment, the method further comprises an act of sharing, by the group of virtual computer systems, the common storage location that stores the plurality of web service requests.

According to one aspect, a system for processing web service transactions is provided. The system comprises a plurality of servers each adapted to receive and process one

or more web service requests, the plurality of servers comprising a first and second server of the plurality of servers that are each configured to compare timestamps associated with at least one database record of a common database associated with a web service application. In one embodiment, the first server is adapted to delete the at least one database record of the common database associated with the web service application, if it is determined that timestamps of the first and second servers have expired, the timestamps being associated with the at least one database record of the common database. In one embodiment, the first and second servers are configured to update a timestamp associated with the at least one database record of the common database associated with the web service application responsive to an access to the at least one database record. In one embodiment, the first and second servers are located in a first and a second datacenter, respectively. In one embodiment, the plurality of servers includes a plurality of virtual servers. In one embodiment, the comparing of timestamps is responsive to a determination by one of the first and second servers that at least one of the timestamps is close to an expiration time. In one embodiment, the system further comprises a request handler element adapted to distribute a web service request to at least one of the plurality of servers for processing. In one embodiment, the system further comprises a component configured to monitor processing of web service requests by the at least one of the plurality of servers. In one embodiment, the system further comprises a component to determine a failover of the processing of at least one web service request to another one of the plurality of servers upon failure of the at least one server. In one embodiment, the web service requests include at least one of a group comprising a REST service request; and a SOAP service request. In one embodiment, the plurality of servers that receive and process web service requests are operable in at least one of a group of operating modes, the group comprising an active/single mode; an active/standby mode; and an active/active mode.

According to one aspect, a method for processing web service transactions is provided. The method comprises acts of receiving and processing, by a plurality of respective servers, one or more web service requests, wherein the method further comprises acts of comparing, by a first and second server of the plurality of servers are to compare timestamps associated with at least one database record of a common database associated with a web service application. In one embodiment, the method further comprises an act of deleting, by the first server, the at least one database record of the common database associated with the web service application, if it is determined that timestamps of the first and second servers have expired, the timestamps being associated with the at least one database record of the common database. In one embodiment, the method further comprises an act of updating, by the first and second servers, a timestamp associated with the at least one database record of the common database associated with the web service application responsive to an access to the at least one database record. In one embodiment, the method further comprises an act of locating the first and second servers in a first and a second datacenter, respectively. In one embodiment, the plurality of servers includes a plurality of virtual servers.

In one embodiment, the act of comparing of timestamps is responsive to an act of determining, by one of the first and second servers, that at least one of the timestamps is close to an expiration time. In one embodiment, the method further comprises an act of distributing, by a request handler element, a web service request to at least one of the plurality

5

of servers for processing. In one embodiment, the method further comprises an act of monitoring processing of web service requests by the at least one of the plurality of servers. In one embodiment, the method further comprises an act of determining a failover of the processing of at least one web service request to another one of the plurality of servers upon failure of the at least one server. In one embodiment, the web service requests include at least one of a group comprising a REST service request; and a SOAP service request. In one embodiment, the plurality of servers that receive and process web service requests are operable in at least one of a group of operating modes, the group comprising an active/single mode; an active/standby mode; and an active/active mode.

Still other aspects, examples, and advantages of these exemplary aspects and examples, are discussed in detail below. Moreover, it is to be understood that both the foregoing information and the following detailed description are merely illustrative examples of various aspects and examples, and are intended to provide an overview or framework for understanding the nature and character of the claimed aspects and examples. Any example disclosed herein may be combined with any other example in any manner consistent with at least one of the objects, aims, and needs disclosed herein, and references to “an example,” “some examples,” “an alternate example,” “various examples,” “one example,” “at least one example,” “this and other examples” or the like are not necessarily mutually exclusive and are intended to indicate that a particular feature, structure, or characteristic described in connection with the example may be included in at least one example. The appearances of such terms herein are not necessarily all referring to the same example. The foregoing and other features and advantages of the invention will be apparent to those of ordinary skill in the art from the following more particular description of the invention and the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram showing a system for caching web service commands according to various aspects of the present invention;

FIGS. 2A-2C show embodiments of various caching modes according to various embodiments of the present invention;

FIG. 3 shows a diagram of different web service requests that may be serviced according to various embodiments of the present invention;

FIG. 4 shows an example configuration of a caching service according to one embodiment of the present invention;

FIGS. 5A-5B show example processes performed by a caching service according to one embodiment of the present invention;

FIG. 6 shows an example computer system with which various aspects of the invention may be practiced;

FIG. 7 shows one example implementation of a caching service using logical partitions distributed among multiple data centers according to one embodiment of the present invention;

FIG. 8 shows a more detailed implementation of a caching service that utilizes multiple caching servers within a logical partition according to one embodiment of the present invention;

6

FIG. 9 shows a process of caching service requests and storing requests in logical storage according to one embodiment of the present invention;

FIG. 10 shows a process for updating timestamp information according to one embodiment of the present invention;

FIG. 11 shows an example computer system upon which various embodiments of the present invention may be practiced;

FIG. 12 shows an example storage system capable of implementing various aspects of the present invention.

FIG. 13 shows a block diagram of an online caching system;

FIG. 14 shows a block diagram of an online caching system;

FIG. 15 shows a block diagram of an online caching system;

FIG. 16 shows a block diagram of an online caching system;

FIG. 17 illustrates a method of processing web service requests;

FIG. 18 illustrates a method of processing web service requests;

FIG. 19 illustrates a method of processing web service requests;

FIG. 20 illustrates a method of processing web service requests;

FIG. 21 illustrates a method of processing web service requests;

FIG. 22 illustrates a first and a second database record stored in a first and a second online storage;

FIG. 23 illustrates a flow diagram for a timestamp exchange process; and

FIG. 24 illustrates a method of processing a caching service request from a client.

DETAILED DESCRIPTION OF EMBODIMENTS OF THE INVENTION

As discussed above, embodiments of the present invention relate to caching systems, and more specifically to a system and method for processing web-based service requests. Disclosed are caching systems capable of handling web service requests from one or more clients.

FIG. 1 shows a block diagram of a distributed system 100 suitable for implementing various aspects of the invention. In particular, according to one embodiment, FIG. 1 shows a system 100 that includes a caching service 102 that is capable of storing and servicing web service commands according to various embodiments of the present invention. More particularly, a caching service 102 may be provided that services one more web service commands received from one more clients such as client 101. Such web service commands may be serviced as part of a website, application service, storage or other web service type. It may be beneficial to provide a storage and command processing platform that is robust and capable of being scaled to any type of application.

FIGS. 2A-2C show various embodiments of different caching service modes according to embodiments of the invention. For example, FIG. 2A shows an active single mode where a single instance of the caching service (e.g., caching service 210) provides support and servicing of client requests (e.g. from client 201A). In what is referred to herein as an active/active configuration and shown in FIG. 2B, a client (e.g., client 201B) sends a request to a load manager 221A that distributes requests to two or more caching

services (e.g. caching service **220A**, caching service **220B**). In this configuration, the load manager sends requests to multiple services, and replication ensures that each caching service is consistent. FIG. **2C** shows an active/standby configuration where one caching service (e.g., caching service **A 230A**) receives and processes requests through an active port, and where one or more other caching services (e.g., caching service **B 230B**) remains “closed” and does not receive and process requests directly.

According to one embodiment of the present invention, it is appreciated that a caching service may be configured using a virtualized environment with logical storage shared between virtual servers, and the caching service may direct web service requests to an available virtual server suitable for servicing the request. For instance, the caching service may direct the request to the most underutilized virtual server that is available. In another embodiment, requests may be stored in a common location accessible by multiple virtual servers, and therefore, replication of request data is not necessary between servers that share the common storage location. Further, upon failure of a particular server, another server may be tasked to service the request, accessing the request in common storage.

FIG. **3** shows a diagram of different web service requests that may be serviced according to various embodiments of the present invention. For instance, a distributed system **300** may be used that includes a client **301** that initiates one or more web service requests that are serviced by a caching service (e.g., caching service **308**). The service requests may include, but are not limited to, HTTPS (e.g., with SSL/TLS encryption) **306**, HTTP requests **305**, secure SOAP (e.g., using web service security (WS-Security or WSS) **304**, SOAP service calls **303**, REST calls **302**, or other types of service requests **307**. It should be appreciated that any type of request or message format may be used with various aspects of the present invention.

FIG. **4** shows an example configuration of a caching service according to one embodiment of the present invention. It is appreciated that a caching service may be operated among multiple data centers in a distributed computer system **400** according to one embodiment of the present invention. For instance, requests received by clients (e.g., client **A 401A**, client **Z 401B**, etc.) may be directed to different data centers (e.g., data center **AA 404A**, data center **ZZ 404B**, data center **BB 404C**) by a load manager (e.g., load manager **402**). Load manager **402** may be a router, a computer system or other type of specialized system capable of receiving requests and directing the requests to different data centers. Each data center may include one or more nodes or systems capable of servicing such requests. Such nodes may include, for example, node **AA 403A**, node **BB 403B**, node **ZZ 403C**. It is appreciated that such nodes may include physical systems, such as one or more server systems, a collection of systems (e.g., bladed server systems), a cluster of systems, mainframe computer system, or any type and configuration of systems. Such systems may also include virtualized resources such as virtual server systems, and may include one or more levels of virtualization. In one embodiment, such a caching system may be implemented across data centers to provide redundancy in the case of site or system failure.

FIGS. **5A-5B** show example processes performed by a caching service according to one embodiment of the present invention. In particular, FIG. **5A** shows an active/active configuration where requests are sent to a number of systems simultaneously. In the event of failure, one of the active caching systems continues to operate. In the example shown,

a client (e.g., client **501**) sends a request to a load manager (e.g., load manager **502**) which directs the request to a caching service. According to one embodiment, the caching service may be a virtual caching service implemented on a virtual computer system. For instance, the virtual computer system may include one or more virtual servers that receive requests and direct them to one or more systems to be processed. For instance, virtual caching service **503** may direct a request in an active/active configuration to both a system **A 504** and a system **B 505** for processing. In an active/active configuration, both system **A** and system **B** receive and process the request and provide an acknowledgment and response to the virtual caching service **503**. The virtual caching service **503** provides a response directly to client **501**.

FIG. **5B** shows an active/standby configuration where two or more caching service entities receive one of the web service requests, but a single entity is responsible for committing transactions. The client **501**, load manager **502**, virtual caching service **503**, and systems (e.g., systems **504**, **505**) may be configured to operate in active/standby or active/active mode as discussed above. According to one embodiment, the virtual caching service **503** directs a request to one of the systems (e.g., system **504**) responsible for committing the transaction, and the responsible system replicates any necessary data to the other system (e.g., system **505**). Virtual caching service **503** is configured to respond to client **501** once a response is received from the responsible system (e.g., system **504**). As the standby system (e.g., system **505**) does not receive the request directly, the standby system takes over only upon failure of the primary system (e.g., system **A**).

FIG. **6** shows an example computer system with which various aspects of the invention may be practiced. For example, a distributed system **600** may be provided that caches web service requests according to various embodiments of the present invention. System **600** includes one or more caching systems (e.g., **601A-601Z**) that operate to receive and process requests. System **600** may also include a load manager **602** that directs receive requests (e.g., requests **603**) to one or more caching systems for processing. Each caching system may also include a request handler **604** that determines, within the caching system, what systems will be processing each request. In one example implementation, a number of servers (e.g., servers **6058-605Z**) are configured to receive requests directed to them from the request handler **604**. Each of the servers may be provided access to a logical storage system **606** where transactions may be committed. Further, requests that are received by the request handler **604** may be stored within a common area whereby servers may access and process received requests. If a failure occurs, and the request is not processed, another server may be capable of servicing the request.

FIG. **7** shows one example implementation of a caching service using logical partitions distributed among multiple data centers according to one embodiment of the present invention. In particular, it should be appreciated that a caching server may be implemented within a system having logical partitions. A logical partition (LPAR) is the division of a computer processor, memory, and storage into multiple sets of resources so that each set of resources can be operated independently with its own operating system instance and applications. A logical partition may be defined that exists on a single computer system or group of computer systems. For example, many systems, including but not

limited to standalone computer systems, groups of computers, clusters, mainframes, and other types of systems may use logical partitions.

As shown in FIG. 7, a distributed system 700 may be provided that includes a number of data centers each having one or more request handlers and logical partitions. For instance, the caching service in a data center #1 (item 701A) may include a request handler (TCP/SD) 703A similar to request handler 604 of FIG. 6. Further, a data center #2 may include a similar request handler (TCP/SD) 703B. Such handlers may receive requests and allocate them to processing entities to be serviced.

As discussed, a caching service may include one or more servers or other processing entities that are capable of processing requests. For instance, as discussed, such entities may include logical partitions or LPARs. In the example shown, the data center #1 may include a number of LPARs (e.g., LPAR #1-LPAR #6 (items 702A-702F)) that are adapted to receive and process requests from request handler 703A. Further, in the example shown, a data center #2 may include a number of LPARs (e.g., LPAR #1-LPAR #3 (items 704A-704C)) that are adapted to receive and process requests from request handler 703B. According to one embodiment, users, organizations, applications or other entities may be assigned to a particular caching service, and that caching service may have an assignment of particular resources (e.g., LPARs, storage, etc.) to fulfill the service requirements for web requests associated with the caching service. Such resources may be allocated to particular entities using, for example, a management interface that is used to set up the caching service for a particular entity.

FIG. 8 shows a more detailed implementation of the caching service that utilizes multiple caching services within a logical partition according to one embodiment of the present invention. For instance, a distributed system 800 is shown including a data center #1 and a defined caching service. System 800 includes a load manager 804 that performs similar functions as discussed above. The caching service also includes a request handler 803 that receives web service requests.

In one implementation, within a particular LPAR (e.g., LPAR #1 (item 802A)), one or more caching servers (e.g., caching servers 805A-805Z) may be defined that are capable of servicing web service requests. Such servers may include, for example, one or more virtual servers that are defined within the particular LPAR. Other LPARs (e.g., LPAR 802Z) may also have caching servers defined within these other LPARs. According to one embodiment, multiple caching servers are permitted to access a logical storage entity for the purpose of servicing requests. Logical storage may include, for example, one or more physical storage devices, servers, or other entities capable of storing data. For instance, a logical storage entity 806 may be commonly accessed by multiple caching servers. The caching servers may share the same logical address space with other servers, and therefore may be capable of servicing requests associated with a particular defined caching service.

In one implementation, the caching service may be implemented in an IBM mainframe environment. For instance, the caching service may be implemented using LPARs as defined in the IBM z/OS environment, as known in the art. Multiple LPARs running z/OS can form a sysplex or parallel sysplex, whether on one machine or spread across multiple machines. Further, such systems may use a logical storage entity such as a VSAM (virtual storage access method) as is used with the z/OS, wherein an enterprise can organize records in a file in physical sequence (the sequential order

that they were entered), logical sequence using a key (for example, the employee ID number), or by the relative record number on direct access storage devices (DASD). Although some examples are shown using an IBM mainframe environment, it should be appreciated that other virtualized computer system types having common storage capabilities may be used, and aspects of the present invention may be used other similarly-configured systems.

FIG. 9 shows an example process of caching service requests and storing requests in logical storage according to one embodiment of the present invention. In a distributed computer system 900, a client 901 communicates a request to a load manager 902. The load manager sends the request to a request handler 903A. In the case of FIG. 9, the caching service is configured in an active/standby mode, where the request is not forwarded to the other request handler of the cache server (e.g., request handler 903B associated with cache server 904B).

In the example shown in FIG. 9, cache server 904A receives the request, performs a post operation on the request which results with a right to logical storage. In the active/standby mode, the first cache server issues a replication post to the request handler of the caching system of data center #2. Request handler 903B sends the request to a suitable cache server (e.g., cache server 904B) which hosts the requests and rights to logical storage (e.g., logical storage 905B). The standby cache server sends a response to the active cache server of data center #1. The active cache server receives the response and sends the response to the client. In this way, the standby and active servers in both data centers are updated, and upon failure of the active server, the standby server may begin servicing requests (e.g., when the load manager 902 forwards requests to a port of the standby server).

In a more detailed example, there are three configurations available to each partition or user of the caching service: active/single; active/standby; or active/active.

In one embodiment, the configuration is set for each partition/user by using a parameter file that specifies the mode of operation. In one embodiment, each active/single data center has its own URUport combination, and clients are configured to referencing the URUport directly in their requests. In an active/standby configuration, the caching system may have a URUport managed by a monitoring component (e.g., a load manager or other entity), and client requests reference the monitoring component (e.g., load manager, load balancer, router or other entity). In the active/standby configuration, a client does not reference the active/standby URUport directly, as only the monitoring component (e.g., a network router and load balancer) references the active/standby URUport.

In one implementation, in the active/active mode, there is a URUport managed by the caching server monitor, with client requests referencing a network router and load balancer. A client does not reference the active/active URUport directly, as only the network router and load balancer references the active/active URUport. When in an active/single configuration, a caching server partition/user is defined to one data center within a virtual computing system. For instance, a z/OS Parallel Sysplex may be used, and as is known, the z/OS Parallel Sysplex combines two basic capabilities of parallel processing and enabling read/write data sharing across multiple systems with full data integrity. A z/OS Parallel Sysplex configuration may include two or more physical zEnterprise servers (CPC/CEC (where CPC is

a central processor complex, and CEC is a central electronic complex or mainframe “box”), two or more logical partitioned operating systems (LPAR) and two or more virtual transaction servers (e.g., implemented as virtual servers, also commonly referred to as “regions”). One virtual transaction server that may be used in the IBM mainframe environment includes the well-known CICS server. The CICS (Customer Information Control System) servers is a well-known family of application servers and connectors provided by IBM that provides industrial-strength, online transaction management and connectivity for mission-critical applications.

In one example implementation, a client application that uses a caching service according to various embodiments sends a REST/SOAP web service request to one URL for an active/single port for a specific data center and the request handler (e.g., a z/OS TCP/SD (Sysplex Distributor) and WLM (work load manager)) routes the request to the ‘best performing’ LPAR and CICS server. Because the request database (e.g., a cache file stored by SMSVSAM and/or CFDT) are accessible to all CICS servers and z/OS LPARs, there is no replication necessary between LPARs and/or CICS regions. With the z/OS Parallel Sysplex, TCP/SD, WLM, multiple LPAR and CICS servers, the active/single provides high availability (HA) and load balancing. Should the entire z/OS Parallel Sysplex fail, there is no recovery or replication required when the system is restarted.

When in an active/standby configuration, a caching server partition/user is defined to one data center within a z/OS Parallel Sysplex as ‘active’ and another data center within a z/OS Parallel Sysplex as ‘standby’. Both active and standby systems are “live” and include two or more zEnterprise servers (CPC/CEC), two or more logical partitioned operating systems (LPAR) and two or more CICS servers (e.g., virtual servers). The client application that uses the caching service sends a REST/SOAP request to one URL, which is handled by a network router and load balancer, which then routes the request to the active/standby port on both data centers. The active system maintains the client port opened, allowing requests from the network router and load balancer. However, the standby system maintains the client port closed. The closed port signals the network router and load balancer to send requests to the active system, which maintains the open client port. The client port on both the active and standby systems are monitored and managed by a caching server background process.

According to one embodiment, while requests are being processed by the active system, asynchronous replication is performed on the request through an internal port only known by caching server to the standby system. The internal port, used by both active/standby and active/active is only known to caching server systems and does not process any requests from caching server clients and/or network router and network load balancers. With the z/OS Parallel Sysplex, TCP/SD, WLM, multiple LPAR and CICS servers, the active/standby provides high availability (HA) and load balancing within the active Sysplex. Should the entire Active z/OS Parallel Sysplex fail, the caching server monitor (background process) on the standby system detects the situation, then immediately opens the client port and sets the caching server monitor control file as ‘in recovery’ status for the system and for each partition/user. When the client port is available, the network router and load balancer then begins routing all active/standby requests to the new ‘active’ system, which previously was marked/designated as ‘standby’. Because caching server requests were being replicated from the other system before the failure, the cached information is readily available when the client port

becomes active, except for messages that were between client response and asynchronous replication. For those records in this situation, the next GET request returns a ‘not found’ status, prompting the client to issue another POST to create the record in the new Active system.

When the failed z/OS Parallel Sysplex, LPARs and the CICS servers are restarted, the active/standby client port is defined as ‘closed’, preventing the network router and load balancer from sending caching server requests to the new ‘Standby’ system, however the internal port known only to caching server is defined as ‘open’. On the active system, a caching server monitor (e.g., a background process) detects that the caching server CICS servers are now available through the internal port, which initiates a background process on the active system side that begins the recovery process. Each caching server partition/user includes a ‘recovery’ task started by the caching server monitor which reads through the caching server file system and issues a POST request across the internal port to the standby system. When each recovery task completes, a status record for each caching server partition/user is updated in the caching server master control file, which is used by the caching server monitor process. During this recovery process, client requests are being processed on the active system, with asynchronous replication being performed on the request through the internal port to the Standby system. Both recovery and replication requests are processed concurrently across the internal port between the active and standby systems. When the recovery task(s) are complete, replication continues for client requests received through the client port on the active system, and the caching server monitor control file is set as ‘recovery complete’ for the system and for each partition/user. The client port on the newly recovered ‘active’ system is opened and the client port on the ‘standby’ system is closed, shifting the workload back to the primary active/standby configuration for those partitions/users within the two data center clusters.

When in the active/active configuration, a partition/user that uses the caching service is defined to two data centers, providing a z/OS Parallel Sysplex each defined as ‘active’. Both active systems are “live” and include two or more zEnterprise servers (CPC/CEC), two or more logical partitioned operation systems (LPAR) and two or more CICS servers (e.g., implemented as virtual servers).

According to one implementation, a client application that uses the caching service sends a REST/SOAP request to one URL, which is handled by a network router and load balancer, which then routes the request to the active/active port on both data centers. The client port on both active/active systems is opened allowing the network router and load balancer to send requests to both active/active systems. While requests are being processed by the active system that receives the request, synchronous replication is being performed on the request through an internal port only known by caching server. The internal port, used by both active/standby and active/active is only known to caching server systems and does not process any requests from caching server clients and/or network router and network load balancers. With the z/OS Parallel Sysplex, TCP/SD, WLM, multiple LPAR and CICS servers, the active/active provides high availability (HA) and load balancing within the active Sysplex.

Should an entire active z/OS Parallel Sysplex fail, the caching server monitor (background process) on the other active system detects the situation, and sets the caching server monitor control file as ‘in recovery’ status for the system and for each partition/user. When the failed z/OS

Parallel Sysplex, LPARs and CICS servers are restarted, the active/active client port is defined as 'closed', preventing the network router and load balancer from sending caching server requests and the caching server monitor control file is set as 'in recovery' status for the system and for each partition/user. The internal port on the recovering system is open during restart. On the opposite active system, a caching server monitor (background process) detects the caching server CICS servers are now available through the internal port, which initiates a background process on the active side that begins the recovery process.

Each caching server partition/user includes a 'recovery' task started by the caching server monitor which reads through the caching server file system and issues a POST request across the internal port to the recovering active system. When each recovery task completes, a status record for each caching server partition/user is updated in the caching server master control file on both systems, which is used by a monitor process of the caching server. During this recovery process, client requests are processed on the active system, with synchronous replication being performed on the request through the internal port to the recovering active system. Both recovery and replication requests are processed concurrently across the internal port between the active and recovering systems. When the recovery task(s) are complete, the port on the recovering system is set to 'opened', enabling the network router and load balancer to send requests to both active/active systems. The caching server monitor control file is set as 'recovery complete' for the system and for each partition/user on both active/active systems.

According to one embodiment, new commands may be provided that implement various basic functions (e.g., POST, GET, PUT and DELETE commands) in a caching system according to various embodiments of the present invention. For example, one process that may be defined according to various embodiments, is a key delete process. For instance, when a DELETE request from the client specifies a key in the URI, a specific record from the caching server is deleted with the request replicated to the opposite (remote) data center. For instance, in a large database, a user may desire to delete a large number of product entries in the database, and thus, by specifying a key range, a single DELETE operation may be performed (e.g., by matching a pattern).

Another option that may be supported by the DELETE request may include a list of keys to be deleted using a 'regex' or regular expression verb on the URI. On a DELETE request with ®ex specified on the URI, caching server will delete a list of keys that match the pattern specified in the ®ex command. Patterns for ®ex may be specified as follows:

/key®ex=*

/key®ex=Sam*

/key®ex=Sam*1

/key®ex=*Sam*

Delete all keys

Delete keys that start with Sam (inclusive of a key <sam')

Delete keys that start with Sam and ending with <1'

Delete keys that contain the word Sam

According to one embodiment, during the delete process, only records with a timestamp of equal or older than 'current' time are deleted. This feature allows records to be added/updated (POST/PUT) during a delete process when ®ex has been specified.

Another feature that may be provided with the delete request when ®ex has been specified includes a 'syn-

chronous delete request' (SDR) or an 'asynchronous delete request' (ADR). The .SDR or .ADR may be specified in the final qualifier of the URI that precedes the key portion of the URI. When .SDR is requested, the delete process is performed synchronously, then the response is returned to the client. When .ADR is requested, the delete process is performed asynchronously after the response has been returned to the client.

In another example, another process that may be defined according to various embodiments, is a key retrieval process or GET. When a GET request from the client specifies a key in the URI, a specific record from caching server is returned. Another option associated with the GET request includes a list of keys to be returned (instead of actual record data) using a 'regex' or regular expression verb on the URI. On a GET request with ®ex specified on the URI, caching server may be adapted to return a list of keys that match the pattern specified in the ®ex command. Patterns for ®ex may be as follows:

/key®ex=*

/key®ex=Sam*

/key®ex=Sam*1

/key®ex=*Sam*

Return all keys

Return keys that start with Sam (inclusive of a key <sam')

Return keys that start with Sam and ending with <1'

Return keys that contain the word Sam

It should be appreciated that other commands or operations may be provided when in a particular caching mode of operation.

In yet another embodiment, a distributed method for updating timestamp information related to stored data is provided. Such timestamp distribution may be used in association with the caching service for web service requests as discussed above. For instance, it is appreciated that timestamp information for data may need to be made consistent across multiple datasets (e.g., located at different datacenters). It is realized that in high-volume transaction systems such as a web-based transaction system, it may not be feasible to adequately transmit timestamp information in an effective manner between systems due to overhead, network traffic, performance, and other considerations.

According to one embodiment, a record expiration process includes an asynchronous or background process that executes on intervals (e.g., as set in a caching server control file) in each of the servers defined in the active/single, active/standby and active/active systems. According to one embodiment, the server startup process starts an expiration task for each partition/user file defined in the servers with the interval defined in the caching server control file. Each expiration task establishes and global ENQ or lock across the Sysplex to ensure serialization of the expiration process for each partition/user. The global ENQ or lock is released when the expiration task completes processing of the partition/user file. Caching server expiration tasks may be started on both systems in the active/standby and active/active configuration.

In active/single, active/standby and active/active configurations, the timestamp on each record gets set on POST/PUT requests when 'last update time' or LUT is specified with a corresponding 'time to live' or TTL value, which may be expressed in seconds. Another option is 'last access time' or LAT where the timestamp on each record gets set on GET/POST/PUT requests with the corresponding TTL value. The minimum TTL may be, for example, 300 seconds. The maximum TTL may be, for example, 86400 seconds. A default value may be set when the TTL value is

not specified. For instance, the default value, when not specified, may be 1800 seconds.

According to one embodiment, one advantages of a caching server over other distributed database products is that records do not get replicated across nodes (e.g., CICS servers) within a cluster (e.g., a Sysplex), as the file systems are accessible to all CICS servers and LPARs within a z/OS Parallel Sysplex. In one implementation, replication of POST/PUT (add/update) requests are performed across active/standby and active/active data centers for all POST/PUT requests.

Another advantage includes, according to one embodiment, how a caching server handles GET and LAT requests, as the timestamp is updated for each of these requests on the local system that receives the request. However, according to one embodiment, these requests are not replicated across data centers. It is appreciated that replicating GET/LAT information across data centers would cause excessive and unnecessary network, processor and VO overhead.

According to one embodiment, a caching server handle does not need to keep timestamps synchronized across data centers when GET/LAT is utilized. Rather, according to one embodiment, timestamp information is used for record expiration, so instead of updating the timestamp on every GET/LAT request, a caching server utilizes a technique called a timestamp exchange during the expiration process. When an expiration task is processing a partition/user file on either the active/standby or active/active systems, each record is read sequentially and the timestamp and time to live (TTL) are used to determine if the record is to be expired. When a record is eligible to be expired on the local system, a request is sent to the opposite (remote) system to delete the expired record. If the record is not eligible to be expired on the remote system, then the timestamp is returned by the timestamp exchange process to the expiration process and the timestamp updated on the local system. If the record is eligible to be expired on the remote system, then the record is deleted on the remote system, and is then deleted on the local system.

Further, a capability may be provided that permits timestamp information to be maintained for a data element based on when the data was last accessed. For instance, most data elements only information that identified when a data element was last updated, not when it was last accessed. To this end, timestamp information may be provided that indicates when the particular data entry was last accessed.

FIG. 10 shows an example process for updating timestamp information according to one embodiment of the present invention. In particular, FIG. 10 shows a process 1000 for determining, among servers distributed among multiple data centers, when a particular database record should be deleted. A block 1001, process 1000 begins. At block 1002, a server (e.g., a virtual server associated with a web request caching service) receives and processes a data request. For instance, there may be one or more operations relate to database entries that may cause the server to update the timestamp associated with a particular database entry. For instance, at block 1003, the server may update a last update timestamp (LUT) associated with a right to a database instance. According to one embodiment, the server may also be capable of updating a timestamp based on a last access (e.g., a last access timestamp (LAT) updated at block 1004.

Further, each of the servers that may be performing operations related to the same database entries may need to determine when such entries should be deleted from the database. For instance, at block 1005, a server determines

whether a timestamp associated with the database entry is close to expiration. If not, the server continues to service database requests. If the timestamp is close to expiration, the server may send a request (e.g., at block 1006) to a corresponding server from another data center (e.g., a server located at data center #2). If it is determined (e.g., at block 1007) that the timestamp associated with that database record is expired, both servers may delete the database record in the databases at both data centers (e.g., datacenters #1 and #2). If not, the datacenter may send its more recent timestamp to data center #1, and the database entry continues to exist. That is, other servers may be operating on database entries, and if another server has a more recent entry, then that database entry should not be deleted. Such timestamps may be checked periodically, eliminating the need for large numbers of messages to make data entries consistent. At block 1010, process 1000 ends.

Example Computer Implementations

Processes described above are merely illustrative embodiments of systems that may be used to cache web service requests. Such illustrative embodiments are not intended to limit the scope of the present invention, as any of numerous other implementations for performing the invention. None of the claims set forth below are intended to be limited to any particular implementation of a caching system, unless such claim includes a limitation explicitly reciting a particular implementation.

Processes and methods associated with various embodiments, acts thereof and various embodiments and variations of these methods and acts, individually or in combination, may be defined by computer-readable signals tangibly embodied on a computer-readable medium, for example, a non-volatile recording medium, an integrated circuit memory element, or a combination thereof. According to one embodiment, the computer-readable medium may be non-transitory in that the computer-executable instructions may be stored permanently or semi-permanently on the medium. Such signals may define instructions, for example, as part of one or more programs, that, as a result of being executed by a computer, instruct the computer to perform one or more of the methods or acts described herein, and/or various embodiments, variations and combinations thereof. Such instructions may be written in any computer programming language available now or in the future.

The computer-readable medium may be transportable such that the instructions stored thereon can be loaded onto any computer system resource to implement the aspects of the present invention discussed herein. In addition, it should be appreciated that the instructions stored on the computer-readable medium, described above, are not limited to instructions embodied as part of an application program running on a host computer. Rather, the instructions may be embodied as any type of computer code (e.g., software or microcode) that can be employed to program a processor to implement the above-discussed aspects of the present invention.

Various embodiments according to the invention may be implemented on one or more computer systems. These computer systems may be, for example, general-purpose computers with any type of processor. Further, the software design system may be located on a single computer or may be distributed among a plurality of computers attached by a communications network.

The computer system may include specially-programmed, special-purpose hardware, for example, an application-specific integrated circuit (ASIC). Aspects of the invention may be implemented in software, hardware or

firmware, or any combination thereof. Further, such methods, acts, systems, system elements and components thereof may be implemented as part of the computer system described above or as an independent component.

A computer system may be a general-purpose computer system that is programmable using a high-level computer programming language. A computer system may also be implemented using specially programmed, special purpose hardware. In a computer system there may be any type of processor known and used now or in the future, with any type of operating system known now or in the future.

Some aspects of the invention may be implemented as distributed application components that may be executed on a number of different types of systems coupled over a computer network. Some components may be located and executed on mobile devices, servers, tablets, or other system types. Other components of a distributed system may also be used, such as databases, cloud services, or other component types.

The processor and operating system together define a computer platform for which application programs in high-level programming languages are written. It should be understood that the invention is not limited to a particular computer system platform, processor, operating system, or network. Further, it should be appreciated that multiple computer platform types may be used in a distributed computer system that implements various aspects of the present invention. Also, it should be apparent to those skilled in the art that the present invention is not limited to a specific programming language or computer system. Further, it should be appreciated that other appropriate programming languages and other appropriate computer systems could also be used.

One or more portions of the computer system may be distributed across one or more computer systems coupled to a communications network. These computer systems also may be general-purpose computer systems. For example, various aspects of the invention may be distributed among one or more computer systems configured to provide a service (e.g., servers) to one or more client computers, or to perform an overall task as part of a distributed system. For example, various aspects of the invention may be performed on a client-server system that includes components distributed among one or more server systems that perform various functions according to various embodiments of the invention. These components may be executable, intermediate (e.g., IL) or interpreted (e.g., Java) code which communicate over a communication network (e.g., the Internet) using a communication protocol (e.g., TCP/IP).

Certain aspects of the present invention may also be implemented on a cloud-based computer system, a distributed computer network including clients and servers, or any combination of systems. It should be appreciated that the invention is not limited to executing on any particular system or group of systems. Also, it should be appreciated that the invention is not limited to any particular distributed architecture, network, or communication protocol.

Various embodiments of the present invention may be programmed using an object-oriented programming language, functional, scripting, and/or logical programming languages. Various aspects of the invention may be implemented in a non-programmed environment (e.g., documents created in HTML, XML or other format that, when viewed in a window of a browser program, render aspects of a graphical-user interface (GUI) or perform other functions).

Various aspects of the invention may be implemented as programmed or non-programmed elements, or any combination thereof.

Further, on each of the one or more computer systems that include one or more components of distributed system **100**, each of the components may reside in one or more locations on the system. For example, different portions of the components of system **100** may reside in different areas of memory (e.g., random access memory (RAM), read-only memory (ROM), disk, etc.) on one or more computer systems. Each of such one or more computer systems may include, among other components, a plurality of known components such as one or more processors, a memory system, a disk storage system, one or more network interfaces, and one or more busses or other internal communication links interconnecting the various components.

Any number of caching systems described herein may be implemented on a computer system described below in relation to FIGS. **11** and **12**. In particular, FIG. **11** shows an example computer system **1100** used to implement various aspects. FIG. **12** shows an example storage system that may be used. System **1100** is merely an illustrative embodiment of a computer system suitable for implementing various aspects of the invention. Such an illustrative embodiment is not intended to limit the scope of the invention, as any of numerous other implementations of the system are possible and are intended to fall within the scope of the invention. For example, a virtual computing platform may be used. None of the claims set forth below are intended to be limited to any particular implementation of the system unless such claim includes a limitation explicitly reciting a particular implementation.

For example, various aspects of the invention may be implemented as specialized software executing in a general-purpose computer system **1100** such as that shown in FIG. **11**. The computer system **1100** may include a processor **1103** connected to one or more memory devices **1104**, such as a disk drive, memory, or other device for storing data. Memory **1104** is typically used for storing programs and data during operation of the computer system **1100**. Components of computer system **1100** may be coupled by an interconnection mechanism **1105**, which may include one or more busses (e.g., between components that are integrated within a same machine) and/or a network (e.g., between components that reside on separate discrete machines). The interconnection mechanism **1105** enables communications (e.g., data, instructions) to be exchanged between system components of system **1100**. Computer system **1100** also includes one or more input devices **1102**, for example, a keyboard, mouse, scanner, trackball, microphone, touch screen, and one or more output devices **1101**, for example, a printing device, display screen, and/or speaker. The system may also include any specialized components depending on the application, including any barcode reader, magnetic stripe reader, receipt printer, hand-held or fixed scanners, pin entry devices (PED), or other device types. In addition, computer system **1100** may contain one or more interfaces (not shown) that connect computer system **1100** to a communication network (in addition or as an alternative to the interconnection mechanism **1105**).

The storage system **1106**, shown in greater detail in FIG. **12**, typically includes a computer readable and writeable nonvolatile recording medium **1201** in which signals are stored that define a program to be executed by the processor or information stored on or in the medium **1201** to be processed by the program. The medium may, for example, be a disk or flash memory. Storage system **1106** may also

include logical storage comprising a number of physical storage elements. Typically, in operation, the processor causes data to be read from the nonvolatile recording medium **1201** into another memory **1202** that allows for faster access to the information by the processor than does the medium **1201**. This memory **1202** is typically a volatile, random access memory such as a dynamic random access memory (DRAM) or static memory (SRAM). It may be located in storage system **1106**, as shown, or in memory system **1104**, not shown. The processor **1103** generally manipulates the data within the integrated circuit memory **1104**, **1202** and then copies the data to the medium **1201** after processing is completed. A variety of mechanisms are known for managing data movement between the medium **1201** and the integrated circuit memory element **1104**, **1202**, and the invention is not limited thereto. The invention is not limited to a particular memory system **1104** or storage system **1106**.

The computer system may include specially-programmed, special-purpose hardware, for example, an application-specific integrated circuit (ASIC). Aspects of the invention may be implemented in software, hardware or firmware, or any combination thereof. Further, such methods, acts, systems, system elements and components thereof may be implemented as part of the computer system described above or as an independent component.

Although computer system **1100** is shown by way of example as one type of computer system upon which various aspects of the invention may be practiced, it should be appreciated that aspects of the invention are not limited to being implemented on the computer system as shown in FIG. **11**. Various aspects of the invention may be practiced on one or more computers having a different architecture or components that that shown in FIG. **11**.

Computer system **1100** may be a general-purpose computer system that is programmable using a high-level computer programming language. Computer system **1100** may be also implemented using specially programmed, special purpose hardware. In computer system **1100**, processor **1103** is any processor used now or in the future that operates any operating system known now or in the future.

FIG. **13** shows an online caching system **1310**. Online caching system **1310** is used to process web-based service requests for one or more clients. Online caching system **1310**, in this embodiment, includes a plurality of servers **1312** coupled to at least one client **1319**. Each one of the plurality of servers **1312** processes caching service requests from client **1319**. Plurality of servers **1312** in this embodiment includes a server **1313**, a server **1314**, a server **1315**, a server **1316**, a server **1317**, and a server **1318**. In some embodiments, there are fewer or more servers included in plurality of servers **1312**. Plurality of servers **1312** is also referred to as a cluster of servers.

Online caching system **1310** also includes an online common storage **1320**. Online common storage **1320** is coupled to each one of plurality of servers **1312**, and is used by each one of plurality of servers **1312** to store key/value data associated with online caching system **1310**. The data portion of the key/value information can be structured or unstructured data, such as text, javascript object notation (JSON), extensible markup language (XML), graphical interchange format (GIF), joint photographic experts group (JPEG), portable document format (PDF), videos, or other format, as the data is treated as binary. When a media or content type of "text/*" is requested, a translation between extended binary coded decimal interchange code (EBCDIC) and American standard code for interchange format (ASCII)

is performed, enabling both *nix and mainframe application to access the same information.

Information is stored in online common storage **1320** in key/value architecture manner, where the key can range from 1 to 255 bytes in size, and the data from 1 byte to 2 gigabytes per record or object. Both the size of the key and the size of the data are determined at the time of an HTTP POST (insert) request. The data record (object) size on a PUT (change) can be different that the original record (object) whether larger or smaller, as the records are variable in length.

Online caching service **1310** is accessible via transmission control protocol/internet protocol (TCP/IP) requests using hypertext transfer protocol (HTTP) and representational state transfer (REST) methodology. Caching system **1310** executes web service requests in one or more of plurality of servers **1312**, where each one of plurality of servers **1312** is primary and none of plurality of servers **1312** require the presence of another one of plurality of servers **1312**. Since each of plurality of servers **1312** is "primary" and does not need any other servers, online caching service **1310** enables the caching service to be executed in any one or more of plurality of servers **1312**, whether first or subsequent requests, to access information stored in online common storage **1320** or other storage devices. Each of plurality of servers **1312** can operate in active/single, active/active, or active standby mode. This provides high availability within plurality of servers **1312** or across a number of plurality of servers, as explained further below.

Each one of plurality of servers (also called nodes) **1312** provides equal functionality independently of each other, and each of plurality of servers **1312** have access to online common storage **1320** as a key/value data architecture. There is no distinguishing between primary and/or secondary servers. Information is stored/retrieved from online common storage **1320** in a shared manner, where online common storage **1320** is accessible to caching services executing in each of plurality of servers **1312**.

Caching service **1310**, when defined in an active/single cluster as shown in FIG. **13**, is located in a single data center location. High availability is achieved because each of plurality of servers **1312** act independently of each other, with each of plurality of servers **1312** sharing online common storage **1320**.

FIG. **14** shows an embodiment of online caching service **1410**. Online caching service **1410** is the same as online caching service **1310** except, in this embodiment, two clients, first client **1319** and a second client **1419** are both coupled to plurality of servers **1312**, and plurality of servers **1312** process caching service requests for each of client **1319** and **1419**. Plurality of servers **1312** can operate in different modes for different clients. In this embodiment, plurality of servers **1312** operates in active/active mode for first client **1312**, and in active/standby mode for second client **1419**.

FIG. **15** shows an embodiment of online caching service **1510**. Online caching system **1510** is the similar to online caching system **1310**, except that in this embodiment an online common storage **1520** is divided among more than one hardware element, and a plurality of servers **1512** is divided among different types of hardware and operating systems. Online caching system **1510** is used to process web-based service requests for one or more clients, in this embodiment a client **1519**. Online caching system **1510** includes a plurality of servers **1512** coupled to at client **1519**. Each one of the plurality of servers **1512** processes caching service requests from client **1519**. Plurality of servers **1512**

in this embodiment includes a server 1513, a server 1514, a server 1515, a server 1516, a server 1517, and a server 1518. In some embodiments, there are fewer or more servers included in plurality of servers 1512. Plurality of servers 1512 is also referred to as a cluster of servers. In this embodiment servers 1513 and 1514 are resident on a first hardware component 1522, where first hardware component 1522 uses a first operating system 1523. Servers 1515 and 1516 are on a second hardware component 1524 with a second operating system 1525, and servers 1517 and 1518 are on a third hardware component 1526 using a third operating system 1527. Plurality of servers 1512 can be divided up onto many different hardware components using many different operating systems.

Online caching system 1510 also includes an online common storage 1520. Online common storage 1520 is coupled to each one of plurality of servers 1512, and is used by each one of plurality of servers 1512 to store key/value data associated with online caching system 1510. In this embodiment, online common storage is resident on a first disk array 1546 and a second disk array 1544. Online common storage can be resident on many different types and numbers of online storage components.

FIG. 16 shows an embodiment of an online caching system 1610. Online caching system 1610 is similar to online caching system 1310 except it is distributed across two site locations, a first site 1628 and a second site 1630. Online caching system 1610 includes a first plurality of servers 1612 resident at first site 1628. First plurality of servers 1612 is coupled to a client 1619. Each one of first plurality of servers 1612 processes caching service requests from client 1619. First plurality of servers 1612 is coupled to first online common storage 1620. Each of first plurality of servers 1612 uses first online common storage 1620 to store key/value data associated with caching system 1610.

Online caching system 1610 also includes second plurality of servers 1632 resident at second site 1630. Second plurality of servers 1632 is communicatively coupled to first plurality of servers 1612 using dark port 1648. Dark port 1648 is used to send communication 1640 between first plurality of servers 1612 and second plurality of servers 1632. A dark port is a port that is used by online caching service 1610, but is not available to client 1619.

Second plurality of servers 1632 is coupled to second online common storage 1642. Each of second plurality of servers 1632 store key/value data associated with online caching system 1610 in second online common storage 1642.

Second plurality of servers 1632 processes caching requests for client 1619, through first plurality of servers 1612, as shown in FIG. 16. The order of processes used for caching service requests from client 1619 using online caching service 1610 depends on whether online caching service 1610 is operating in an active/standby mode or an active/active mode for client 1619. Each of these is explained herein.

When online caching service 1610 is operating in an active/standby mode for client 1619, the caching process begins with client 1619 sending a service request 1650 to first plurality of servers 1612. Service request 1650 can be, for example, a PUT or a POST command, or any other caching service request. In the embodiment shown in FIG. 16, service request 1650 is sent to server 1618, but service request 1650 could be sent to any one of first plurality of servers 1612. First plurality of servers 1612, in this embodiment server 1618, commits service request 1650 to online common storage 1620 in a communication 1658 as shown in

FIG. 16. Since first plurality of servers 1612 is operating in active/standby mode for client 1619, first plurality of servers 1612 next returns success response 1652 to client 1619. Once success response 1652 is sent to client 1619, then first plurality of servers 1612 sends update 1654 to second plurality of servers 1632, requesting second plurality of servers 1632 commit service request 1650 to second online common storage 1642. Second plurality of servers 1632 commits service request 1650 to second online common storage 1642 in communication 1660. Second plurality of servers 1632 sends caching success response 1656 to first plurality of servers 1612, ending the process. In this process, first plurality of servers 1612 sends success report 1652 to client 1619 after caching service request 1650 from client 1619 in first online common storage 1620, and before replicating service request 1650 to second plurality of servers 1632.

When online caching service 1610 is operating in an active/active mode for client 1619, the caching process begins with client 1619 sending service request 1650 to first plurality of servers 1612. First plurality of servers 1612, in this embodiment server 1618, commits service request 1650 to online common storage 1620 in communication 1658 as shown in FIG. 16. Since first plurality of servers 1612 is operating in active/active mode for client 1619, first plurality of servers 1612 next sends update 1654 to second plurality of servers 1632, requesting second plurality of servers 1632 commit service request 1650 to second online common storage 1642. Second plurality of servers 1632 commits service request 1650 to second online common storage 1642 in communication 1660. Second plurality of servers 1632 sends caching success response 1656 to first plurality of servers 1612. After second plurality of servers 1632 replicates the caching of service request 1650, first plurality of servers 1612 sends success response 1652 to client 1619, ending the process. In this process, first plurality of servers 1612 sends success report 1652 to client 1619 after caching service request 1650 from client 1619 in first online common storage 1620, and after second plurality of servers 1632 replicates service request 1650 to second online common storage 1642.

FIG. 17 and FIG. 18 illustrate the methods of processing web service requests as described above using online caching system 1610 of FIG. 16.

FIG. 17 illustrates method 1700 of processing web service requests. Method 1700 includes an act 1710 of caching a service request from a client in a first online common storage at a first site, where the first online common storage is coupled to a first plurality of servers. Method 1700 also includes an act 1720 of replicating the service request using a dark port to a second plurality of servers at a second site. In some embodiments, the dark port is a synchronous dark port. And method 1700 includes an act 1730 of caching the service request in a second online common storage, where the second online common storage is coupled to the second plurality of servers. And method 1700 includes an act 1740 of delivering a success response to the client. Method 1700 can include many other acts.

In some embodiments, method 1700 includes the second plurality of servers sending a response back to the first plurality of servers that the service request has been cached in the second online common storage, before delivering the success response to the client. In some embodiments, delivering the success response to the client occurs after caching the service request from the client in the first online common storage at the first site, and before replicating the service

request using the dark port to the second plurality of servers at the second site. In some embodiments, the dark port is an asynchronous dark port.

In some embodiments, method **1700** includes the second plurality of servers sending a response back to the first plurality of servers that the service request has been cached in the second online common storage after caching the service request in a second online common storage.

FIG. **18** illustrates a method **1800** of processing web service requests. Method **1800** includes an act **1810** of caching a service request from a client in a first online common storage at a first site, where the first online common storage is coupled to a first plurality of servers. And method **1810** includes an act **1820** of delivering a success response to the client. Method **1800** also includes an act **1830** of replicating the service request using a dark port to a second plurality of servers at a second site. And method **1800** includes an act **1840** of caching the service request in a second online common storage, where the second online common storage is coupled to the second plurality of servers. In some embodiments, method **1800** includes the second plurality of servers sending a response back to the first plurality of servers that the service request has been cached in the second online common storage after caching the service request in a second online common storage.

FIG. **19** illustrates a method **1900** of processing web service requests using any one of online caching systems **1310**, **1410**, **1510**, or **1610**. Method **1900** includes an act **1910** of receiving a POST command from a client, where the POST command is associated with a POST request key and a POST request value. Method **1900** also includes an act **1920** of checking whether the POST request key exists in a cache memory with an associated stored value. And method **1900** includes an act **1930** of replacing the stored value with the POST request value in response to determining that the POST request key value already exists in the cache memory.

Method **1900** can include many other acts. In some embodiments, method **1900** includes receiving a PUT command from a client, wherein the PUT command is associated with a PUT request key and a PUT request value. In some embodiments, method **1900** includes inserting the PUT request value into a cache memory in response to determining that the PUT request key does not exist in the cache memory. In some embodiments, method **1900** includes checking whether the PUT request key exists in a cache memory. In some embodiments, method **1900** includes inserting the PUT request value into the cache memory in response to determining that the PUT request key does exist in the cache memory, wherein the PUT request value has a memory size different from the memory size of the value originally associated with the PUT request key.

In some embodiments, method **1900** includes receiving a DELETE command from a client, wherein the DELETE command is paired with a CLEAR querystring and a DELETE timestamp, and wherein the DELETE command is not associated with a key. In some embodiments, method **1900** includes deleting each key/value pair stored in the cache memory that has a timestamp older than the DELETE timestamp.

In some embodiments, method **1900** includes performing a translation between EBCDIC and ASCII in response to receiving a value with a text content type. In some embodiments, method **1900** includes determining a size of the POST request key and a size of the POST request value data in response to receiving the POST command.

FIG. **20** illustrates a method **2000** of processing a web service request that includes an act **2010** of receiving a PUT

command from a client, where the PUT command is associated with a PUT request key and a PUT request value. Method **2000** also includes an act **2020** of inserting the PUT request value into a cache memory in response to determining that the PUT request key does not exist in the cache memory. Method **2000** can include many other acts.

FIG. **21** illustrates a method **2100** of processing a web service request. Method **2100** includes an act **2110** of receiving a DELETE command from a client, where the DELETE command is paired with a CLEAR querystring and a DELETE timestamp, and where the DELETE command is not associated with a key. Method **2100** also includes an act **2120** of deleting each key/value pair stored in a cache memory that has a timestamp older than the DELETE timestamp. Method **2100** can include many other acts.

FIG. **16**, FIG. **22**, FIG. **23**, and FIG. **24** illustrate a disclosed timestamp exchange process executed by online caching system **1610**. FIG. **16** illustrates a simplified block diagram of online caching system **1610**. FIG. **22** illustrates a simplified diagram of a plurality of database records stored in online storage **1620** and **1642** by online caching system **1610**. FIG. **23** illustrates a flowchart of a timestamp exchange process **2300** executed by online caching system **1610**. FIG. **24** illustrates a method **2400** of processing a caching service request from a client conducted by online caching system **1610**.

Referring to FIG. **16** and FIG. **22**, online caching system **1610** includes first server **1618** resident at first site **1628**. First server **1618** is one of first plurality of servers **1612**. First server **1618** is communicatively coupled to client **1619** using a communication link **1651**. First server **1618** processes caching service requests for client **1619**, such as a caching service request **1650** as shown in FIG. **16**. First server **1618** caches database records in first online storage **1620** as part of processing caching service requests for client **1619**. For example, first server **1618** stores a first database record **2120** (FIG. **22**) in first online storage **1620** as part of processing caching service request **1650**. First online storage **1620** is resident at site **1628** and is communicatively coupled to server **1618** using a communication link **1653**.

Each caching instance (also called a service) has its own storage area, whether in memory, on disk or a combination of both. An 'entry' in the cache storage area is represented as a record and accessed as a key/value. Each record will have a timestamp (expressed as absolute time of when the record was created/updated), a time-to-live (TTL, represented in seconds), media-type (aka content-type) and the number of 32 kb segments.

First server **1618** stores first database record **2120** in first online storage **1620**, as shown in FIG. **16** and FIG. **22**. First database record **2120** includes key/value data **2130**, as well as a first timestamp **2150**, and a first TTL **2160**, as well as other elements. Key/value data **2130** is associated with caching service request **1650**, in this embodiment.

It is to be understood that first server **1618** caches a plurality of database records in online storage **1620**, first database record **2120** is one of these records. First server **1618** caches a database records **2122**, for example, which includes key/value data **2132**, timestamp **2152**, and TTL **2162**, as well as any number of further database records, as illustrated in FIG. **22**.

Online caching system **1610** also includes second server **1635** resident at second site **1630**. Second server **1635** is one of second plurality of servers **1632**. Second server **1635** is communicatively coupled to first server **1618** using a communication link **1648**, which is a dark port, in this embodiment. Second server **1635** processes caching service

25

requests for client 1619 through first server 1618, in this embodiment, such as caching service request 1650 as shown in FIG. 16. Second server 1635 caches database records in second online storage 1642 as part of processing caching service requests for client 1619. Second server 1635 is communicatively coupled to second online storage 1642 using a communication link 1661. Second server 1635 stores a second database record 2121 in second online storage 1642 (FIG. 16 and FIG. 22) in response to receiving caching service request 1650 from first server 1618.

Second server 1635 stores second database record 2121 in second online storage 1642, as shown in FIG. 22. Second database record 2121 includes key/value data 2130, as well as a second timestamp 2151 and a second TTL 2161. Key/value data 2130 is associated with caching service request 1650.

It is to be understood that second server 1635 caches a plurality of database records in second online storage 1642, second database record 2121 is one of these records. Second server 1635 caches a database records 2123, for example, which includes key/value data 2133, timestamp 2153, and TTL 2163, as well as any number of further database records, as illustrated in FIG. 22.

First database record 2120 and second database record 2121 each have the same key/value data 2130, which is associated with caching service request 1650. First database record 2120 and second database record 2121 are stored at different sites and in different storage media, and have different timestamps and TTL values. First and second database records 2120 and 2121 are stored temporarily, and are to be deleted when their lifetime has been exceeded. But before they are deleted, a check is done to see if another version of the key/value data has a more updated timestamp and/or TTL. If so, the database record timestamp is updated instead of deleting the database record. Online caching system 1610 using a timestamp exchange process to update and/or delete records. FIG. 23 shows a simplified flowchart of this timestamp exchange process.

The timestamp exchange process 2300 runs on 30 minute intervals on server 1618, in this example, and reads thru the cache storage sequentially using the key to retrieve the value (record) and checks to see if a timestamp has expired. A timestamp is expired when the associated TTL value has been exceeded, which is determined by subtracting the timestamp value stored in the record entry from the current absolute time (the time that this process is being performed), and comparing the resulting difference to the TTL. For example, and as shown in FIG. 23, at node 2320 first database record 2120 is retrieved, and at node 2330 the timestamp 2150 and TTL 2160 of first database record 2120 are retrieved. At node 2340, timestamp 2150 and TTL 2160 of first database record 2120 are used to determine if database record 2120 is eligible for expiration, meaning the life of first database record 2120 has expired and database record 2120 is to be deleted. Server 1618 determines if database record 2120 is eligible for expiration by checking to see if first timestamp 2150 has expired. First timestamp 2150 has expired when first TTL 2160 is exceeded. Server 1618 determines if first TTL 2160 has been exceeded by subtracting first timestamp 2150 from the current absolute time and determining if the difference is larger than the TTL value, which would mean first timestamp 2150 has expired, first TTL 2160 has been exceeded, and first database record 2120 has exceeded its time-to-live 2160.

If first database record 2120 is not eligible for expiration, then the process returns to node 2320 and retrieves the next record.

26

If first database record 2120 is eligible for expiration (if timestamp 2150 has expired because TTL 2160 has been exceeded), then at node 2342, a call is made to a secondary data center using the current record key, which, in this example, is second site 1630 and key/value 2130. In the embodiment shown in FIG. 16, node 2342 of FIG. 23 is executed by first server 1618 sending a timestamp update request 1653 to second server 1635 regarding key/value data 2130. For node 2344 of FIG. 23, second server 1635 retrieves second database record 2121, which shares key/value data 2130 with first database record 2120. At node 2350, second server 1635 uses second timestamp 2151 and second TTL 2161 to determine if second database record 2121 is eligible for expiration. Second database record 2121 is eligible for expiration when second timestamp 2151 is expired. Second timestamp 2151 is expired when second TTL 2161 has been exceeded. Second TTL 2161 has been exceeded when the difference between second TTL 2161 and the current absolute time is larger than second TTL 2161. At node 2350, second timestamp 2151 is subtracted from the current absolute time to determine if second TTL 2161 has been exceeded. If second TTL 2161 has been exceeded, second timestamp 2151 is expired. If second TTL 2161 has not been exceeded, second timestamp 2151 has not expired.

If second TTL 2161 has been exceeded, meaning database record 2121 is eligible for expiration, several actions occur (see FIG. 23): First, at node 2360, second database record 2121 is deleted from second online storage 1642. Second, at node 2362, a timestamp expired indication 1655 is sent from second server 1635 to first server 1618. Timestamp expired indicator 1655 tells first server 1618 that second database record 2121 has also expired (second timestamp 2151 has expired and second TTL 2161 is exceeded), and that both first and second database record 2121 and 2120 are to be deleted. Third, at node 2364, first server 1618 deletes first database record 2120 in response to receiving timestamp expired indicator 1655 from second server 1635.

However, if second TTL 2161 has not been exceeded, meaning second database record 2121 is not eligible for expiration (second timestamp 2151 has not expired and second TTL 2161 has not been exceeded), several different actions occur (see FIG. 23): First, second database record 2121 is not deleted from second online storage 1642, but instead, at node 2370, second server 1635 sends second timestamp 2151 back to first server 1618. Second, at node 2372, first server 1618 replaces first timestamp 2150 with second timestamp 2151, and both first and second database records 2120 and 2121 remain in first and second online storage 1620 and 1642.

This timestamp exchange process 2300 eliminates the need to be sending the whole database record 2120 or key/value data 2130 back and forth to synchronize data between sites 1628 and 1630. Timestamp exchange process 2300 has been developed to expire a record (entry) or sync the timestamps between primary and secondary data centers such as site 1628 and site 1630. Timestamp exchange process 2300 has significantly reduced the CPU and network overhead involved in keeping the timestamps in sync between primary and secondary data centers for online caching system 1610. Timestamp exchange process 2300 eliminates 'thrashing' by only syncing the timestamp when one side or the other expires, and by deleting when both sides have expired.

Thus, using timestamp exchange process 2300, first server 1618 deletes database record 2120 in response to determining that both a) first timestamp 2150 of first database record 2120 and stored in first online storage 1620 has expired, and

b) second timestamp **2151** of second database record **2121** and stored in second online storage **1642** has expired. First timestamp **2150** is expired when first timestamp **2150** indicates that first TTL **2160** has expired—been exceeded. First timestamp **2150** indicates that first TTL **2160** has been exceeded when the difference between first timestamp **2150** and the current absolute time is larger than first TTL **2160**. When the different is exceeded, first database record **2120** has expired—exceeded its time to live. Second timestamp **2131** is expired when second timestamp **2151** indicates that second TTL **2161** has expired—been exceeded. When second TTL **2161** has been exceeded, second database record **2121** has expired—exceeded its time to live.

If both first and second timestamps **2150** and **2151** has expired, both first and second database records **2120** and **2121** are deleted. If first timestamp **2150** is expired but second timestamp **2151** has not expired, first timestamp **2150** is replaced with second timestamp **2151**, and both database records **2120** and **2121** are kept.

FIG. **24** illustrates a method **2400** of processing a caching service request from a client. FIG. **24** includes an act **2410** of caching a first database record in a first online storage. The first database record includes a key/value data associated with the caching service request and a first timestamp.

Method **2400** also includes an act **2420** of caching a second database record in a second online storage. The second database record is also associated with the caching service request, and includes the key/value data and a second timestamp. In some embodiments, the first online storage is located at a first site, and the second online storage in located at a second site at a location different than the first site.

Method **2400** also includes an act **2430** of replacing, by a first server communicatively coupled to the first online storage, the first timestamp with the second timestamp in response to determining that the first timestamp has expired and the second timestamp has not expired.

In some embodiments, method **2400** includes sending, by the first server, the caching service request to the second server.

In some embodiments, method **2400** includes determining, by the first server, that the first timestamp has expired. In some embodiments, determining that the first timestamp has expired includes subtracting the first timestamp from the current absolute time and comparing the resulting value to a first time-to-live (TTL) value. The first TTL value is associated with the first database record. If the difference between the first timestamp and the current absolute time is greater than the first TTL value, the first TTL value has been exceeded and the first timestamp has expired. If the difference between the first timestamp and the current absolute time is less than the first TTL value, the first timestamp has not expired.

In some embodiments, method **2400** includes determining, by the second server, that the second timestamp has expired. In some embodiments, method **2400** includes determining, by the second server, that the second timestamp has not expired. In some embodiments, determining whether the second timestamp has expired includes subtracting the second timestamp from the current absolute time and comparing the resulting value to a second time-to-live (TTL) value. The second TTL value is associated with the second database record. If the difference between the second timestamp and the current absolute time is greater than the second TTL value, the second TTL value has been exceeded, and the second timestamp has expired. If the difference between the second timestamp and the current absolute time is less than

the second TTL value, the second TTL value has not been exceeded, and the second timestamp has not expired.

In some embodiments, method **2400** includes sending a timestamp update request from the first server to the second server. In some embodiments, method **2400** includes sending, by the second server, the second timestamp to the first server. If the first timestamp has expired but the second timestamp has not expired, the second server sends the second timestamp to the first server, and the first server replaces the first timestamp with the second timestamp (act **2430**).

In some embodiments, method **2400** includes deleting the first database record and the second database record in response to determining that both the first timestamp and the second timestamp have expired. If both the first and the second timestamps are expired, first and the second database records are deleted. In some embodiments, method **2400** includes receiving, by the first server, an indication that the second timestamp has expired from the second server. If the second server determines that the second timestamp has expired, the second server sends the indication that the second timestamp has expired to the first server.

The embodiments and examples set forth herein were presented in order to best explain the present invention and its practical application and to thereby enable those of ordinary skill in the art to make and use the invention. However, those of ordinary skill in the art will recognize that the foregoing description and examples have been presented for the purposes of illustration and example only. The description as set forth is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the teachings above.

The invention claimed is:

1. An online caching system comprising:

a first server resident at a first site, wherein the first server is communicatively coupled to a client, and wherein the first server processes a caching service request from the client;

a first online storage communicatively coupled to the first server, wherein the first server stores a first database record comprising a key/value data in the first online storage;

a second server resident at a second site, wherein the second server is communicatively coupled to the first server by a dark port, wherein the dark port is not available to the client for communication; and

a second online storage communicatively coupled to the second server, wherein the second server stores a second database record comprising the key/value data in the second online storage;

wherein the first server deletes the first database record in response to determining that both a) a first timestamp of the first database record and stored in the first online storage has expired, and b) a second timestamp of the second database record and stored in the second online storage has expired.

2. The online caching system of claim **1**, wherein the second server deletes the second database record in response to determining that both a) the first timestamp has expired, and b) the second timestamp has expired.

3. The online caching system of claim **1**, wherein the first server and the second server are each operable in either an active/active mode, or an active/standby mode.

4. The online caching system of claim **1**, wherein the key/value data is associated with the caching service request.

5. The online caching system of claim 1, wherein the second server stores the second database record in the second online storage in response to receiving the caching service request from the first server.

6. A method of processing a caching service request from a client, the method comprising:

5 caching a first database record in a first online storage, wherein the first database record comprises a key/value data associated with the caching service request and a first timestamp; 10

caching a second database record in a second online storage, wherein the second database record comprises the key/value data and a second timestamp;

determining, by the first server, that the first timestamp has expired; 15

sending a timestamp update request from the first server to the second server via a dark port, wherein the dark port is not available to the client for communication;

receiving, by the first server, an indication that the second timestamp has expired; and 20

deleting the first database record and the second database record in response to determining that both the first timestamp and the second timestamp have expired.

7. The method of claim 6, further comprising sending, by the first server, the caching service request to the second server. 25

* * * * *